Maintaining Matroid Intersections Online

Niv Buchbinder* Anupam Gupta† Daniel Hathcock‡ Anna R. Karlin §
Sherry Sarkar¶

Abstract

Maintaining a maximum bipartite matching online while minimizing augmentations is a well studied problem, motivated by content delivery, job scheduling, and hashing. A breakthrough result of Bernstein, Holm, and Rotenberg (*SODA 2018*) resolved this problem up to a logarithmic factors. However, to model other problems in scheduling and resource allocation, we may need a richer class of combinatorial constraints (e.g., matroid constraints).

We consider the problem of maintaining a maximum independent set of an arbitrary matroid $\mathcal{M}$ and a partition matroid $\mathcal{P}$. Specifically, at each timestep t one part P_t of the partition matroid is revealed: we must now select at most one newly-revealed element, but may exchange some previously selected elements, to maintain a maximum independent set on the elements seen thus far. The goal is to minimize the number of augmentations. If $\mathcal{M}$ is also a partition matroid, we recover the problem of maintaining a maximum bipartite matching online with recourse as a special case.

Our main result is an $O(n \log^2 n)$-competitive algorithm, where n is the rank of the largest common base; this matches the current best quantitative bound for the bipartite matching special case. Our result builds substantively on the result of Bernstein, Holm, and Rotenberg: a key contribution of our work is to make use of market equilibria and prices in submodular utility allocation markets.

*Dept. of Statistics and Operations Research, Tel Aviv University, Israel. Supported in part by Israel Science Foundation grant 2233/19 and United States - Israel Binational Science Foundation grant 2022418.

†Carnegie Mellon University and Google Research. This research was partly supported by NSF awards CCF-1955785, CCF-2006953, and CCF-2224718.

‡Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh PA 15213. Supported by the NSF Graduate Research Fellowship grant DGE-2140739.

§University of Washington, Seattle, WA 98195. This research was supported in part by NSF award CCF-1813135 and Air Force Office of Scientific Research grant FA9550-20-1-0212.

¶Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh PA 15213.

1 Introduction

In the *Online Matroid Intersection Maintenance Problem with recourse*, we want to maintain a maximum independent set in the intersection of an arbitrary matroid $\mathcal{M}$ and a partition matroid $\mathcal{P}$ in the online setting. Specifically, suppose the partition matroid is given by a partition $(P_1, P_2, \ldots, P_\ell)$ of the element set E, and $\mathcal{M}$ is another matroid on E. [1] Both of these matroids are initially unknown to us. Now at each timestep t, we have a current maximum independent set I_{t-1} and the next part P_t of the partition matroid is revealed.

Since we need to maintain a maximum independent set in the intersection of $\mathcal{M}$ and the portion of the partition matroid seen so far, we may need to perform some augmentations—i.e., we may need to drop elements from I_{t-1} and add elements from $E \setminus I_{t-1}$ to the current independent set. Our objective is to minimize the total number of reassignments (i.e., additions or deletions from the independent set) over the course of the arrival of all parts of $\mathcal{P}$.

A special case of our problem that has been considered extensively [GKKV95, BLSZ14, BLSZ18, BLSZ22, BHR18] is that of the online *bipartite matching* problem with recourse. This setting corresponds to the matroid $\mathcal{M}$ also being a partition matroid. In turn, it allows us to identify the elements with edges of a bipartite graph whose vertices are the parts in the two partitions. Hence each timestep corresponds to a new vertex from one side of the graph arriving, along with its incident edges.

In order to maintain a maximum matching, we need to augment along alternating paths, which corresponds to dropping or adding edges. (If individual edges arrive one-by-one rather than vertices, nothing better than $\Omega(n^2)$ total cost is possible in the worst case. For example, in the instance where edges arrive on alternating ends of a path, the augmentation must be the entire path at each step [BHR18, §1].)

In addition to its natural and combinatorial appeal, the generality of the online matroid intersection maintenance problem allows us to model problems in resource allocation and scheduling beyond the matching case:

- *Laminar matroids* generalize the bipartite matching setting to allowing constraints on a hierarchy of "groups": e.g., suppose clients arrive online and need to be matched to a server from their desired subset. However, we may have restrictions on the number of clients assigned to servers on the same server rack, or the same data center (due to cooling, power, and bandwidth constraints). These can be captured by laminar matroids, where we are given capacities on a family of laminar sets. Laminar restrictions are common when considering such job scheduling problems with restricted (hierarchical) resources.

- A different setting is that of *matroid partitioning* [Edm65]: the elements of a single matroid $\mathcal{M}$ arrive over time, and need to be partitioned among k color classes, so that each color class is an independent set in $\mathcal{M}$. The goal is to minimize the number of color changes. In this setting, the underlying matroid constraint captures the scheduling constraints of a single server cluster (e.g., like in the laminar case above, or the examples below), the coloring captures the idea of partitioning the jobs among these clusters, and the recourse bound ensures that only a few jobs are reassigned between clusters. The matroid coloring problem can be modeled using online matroid intersection by the simple idea of "lifting" elements to (element, color) pairs.

- *Transversal matroids* are a useful matroid constraint for modern schedulers, since they can model *coflows* [CS12] (which are tasks that can be jointly processed on a computing resource, e.g., can be shuffled in parallel on a MapReduce cluster [IMPP19]). In transversal matroids the elements are *nodes* on one side of a bipartite graph, and independent sets correspond to matchable subsets of nodes. Combining these with the matroid partitioning idea allows us to partition a collection of jobs among clusters. In particular, if we now have several different computing resources (clusters), and jobs arriving online, the scheduler algorithm needs to decide which cluster to choose for each job in order to process it, and the goal is to minimize the number of jobs that have to be switched between clusters.

- If we have a routing problem instead of a scheduling one [JKR17], we can use *gammoids* instead of transversal matroids: these capture sets of nodes which admit vertex-disjoint flows to a sink; again the goal would be to minimize reroutings.

[1] A *matroid* $\mathcal{M}$ over a ground set E is given by a downward-closed collection of independent sets $\mathcal{I} \subseteq 2^E$ such that for any $A, B \in \mathcal{I}$ with $|A| < |B|$ there exists an element $e \in B \setminus A$ such that $A \cup \{e\} \in \mathcal{I}$ [Sch03, §39].

1.1 Our Result and Techniques

The augmenting-path algorithm for matroid intersection [AD71, Law75] (see also [Sch03, §41.2]) immediately bounds the total number of reassignments by $O(n^2)$, where n is the rank of the maximum independent set in $\mathcal{P} \cap \mathcal{M}$. To our knowledge, nothing better was known about this problem in general prior to the current paper. Our main result is the following:

THEOREM 1.1. (MAIN THEOREM) *The* Shortest Augmenting Path *(SAP) algorithm for the* Online Matroid Intersection Maintenance *problem results in at most $O(n \log^2 n)$ total reassignments, where n is the rank of the intersection $\mathcal{M} \cap \mathcal{P}$.*

Our analysis reveals a perhaps surprising connection between matroid intersection maintenance and the well-known theory of market equilibrium prices. Indeed, we define a market with the arriving parts viewed as buyers who want to get one item/element from their desired subset. Viewing the items as being divisible allows us to find a market equilibrium, where the price of each item gives us crucial information about the length of the shortest augmenting path starting with that item. Note that the algorithm is purely combinatorial; the market helps us expose the properties of the underlying instance, and to argue about the algorithm's performance.

The bound of $O(n \log^2 n)$ matches that given by the breakthrough paper of Bernstein, Holm and Rotenberg [BHR18] for the special case of online matchings; this is not a coincidence. Indeed, our first step is to reinterpret their work in the language of market equilibria; we then develop the machinery and connections to reason about general matroid intersection markets and obtain our results. We now elaborate on these connections and our techniques.

Bipartite Matching In the special case of online bipartite matching, we imagine the vertices ("clients") of one side of a bipartite graph, along with the induced edges to the other side ("servers"), arriving online. [BHR18] define a notion of *server necessity* to capture how much of each server s is needed to match all clients. They compute server necessities via "balanced flows" (which is the solution to a certain convex program), and also via an intuitive combinatorial decomposition (also called a "matching skeleton" in other works) that builds on Hall's Theorem. These ingredients give an *expansion lemma* bounding the length of the shortest augmenting path for a new client in terms of the minimum server necessity among its neighbors.

We begin our investigation by showing in §2 how the language and machinery of market equilibria yield a concise and appealing version—though unchanged in its fundamentals—of Bernstein et al's proof for the online bipartite matching problem with recourse. Our starting point is a reinterpretation of the concept of server necessity as *prices in a market equilibrium*. A Fisher market [BS05] consists of n buyers and m divisible items: each buyer i arrives with a budget of money m_i and a utility function that specifies buyers' utilities for each possible bundle of goods. A *market equilibrium* is a set of prices $p_1, \dots, p_m$, where p_j is the price of item j, such that each buyer spends their money on a utility-maximizing bundle, the supply precisely equals the demand, and the market clears (i.e., each good with positive price is sold and each buyer spends all their money). If we view clients in the online matching problem as buyers each having one dollar, and the servers as items, with each buyer having equal and linear utility for each item in their bipartite graph neighborhood (and zero utility for non-neighbors), the market clearing prices turn out to be precisely the server necessities. The equilibrium allocation and prices can be computed using the Eisenberg-Gale convex program [EG59] (which differs from the convex program used in [BHR18]).

Matroid Intersection We then show how this market equilibrium perspective allows us to generalize to online *matroid intersection* with recourse, where we have a more general set of feasibility conditions on the allocations. Again, we start with n buyers and m items, but the items are now the elements E of matroid $\mathcal{M}$, and the buyers are interested in disjoint elements. (In the online matching problem, these elements of the matroid are edges incident to the buyer/online vertex, and hence map to a set of offline vertices that buyer wants to match to.) Each buyer again arrives with some money m_i and a linear utility function over the items in its part. In this market, the allocation of items to buyers must lie in the matroid polytope of $\mathcal{M}$. This market again is an Eisenberg-Gale market [JV07], for which a market equilibrium exists and can be be computed with a

convex program of the form:

$$\max \left\{ \sum_i m_i \log \sum_{e \in P_i} y_e \;\middle|\; \sum_{e \in S} y_e \le \operatorname{rank}_{\mathcal{M}}(S) \quad \forall S \subseteq E, y \ge 0 \right\}.$$

The dual variables to the program again yield equilibrium prices for each item.

Using these ideas, we extend the ideas from §2 to the general online matroid intersection problem in §3.

Our setting requires new ideas beyond the case of matchings because the convex program is richer: the prices (i.e., duals) are now on sets and not on elements. There is a natural way to translate from sets to elements: the price of each element is the sum of prices of sets containing it—but then the prices are not unique, and we can no longer argue the monotonicity of prices as new clients arrive, a crucial ingredient in [BHR18] and in §2. To address this, we first make a connection to submodular utility allocation markets [JV10] to show that prices seen by buyers are monotone. Then we show a decomposition theorem for matroids (extending such a result for matchings [GKK12, BHR18]) that allows us to define unique and consistent prices for elements, and to show monotonicity of all individual element prices over arrivals.

In conjunction with this, we can define for any element e a collection of nested sets, showing that if there are no short augmenting paths (in the natural exchange graph) starting at this element, then these nested sets grow exponentially at rate $\approx (1/p_e)$. With this, we bound the length of paths by $\approx \frac{\ln n}{1-p_e}$ (of course, the paths are never of length more than n.) Finally, the monotonicity of element prices allows us to distribute this augmentation cost to the price increase of the elements participating in this augmenting path. Since the element prices lie in $[0,1]$, the cost charged to each element (over the entire run of the algorithm) is

$$\int_{p=0}^{1} \min\left\{n, \frac{\ln n}{1-p}\right\} dp \leq \int_{p=0}^{1-1/n} \frac{\ln n}{1-p}\, dp + \int_{p=1-1/n}^{1} n\, dp = O(\ln^2 n).$$

A technical detail is that summing this over all elements would give us $|E| \log^2 n$, and not something that depends on the rank. This issue can be handled using a convexity-based argument.

Paper Outline In the remainder of the paper, we first illustrate the basic ideas of our arguments in §2 for the setting of bipartite matchings. In §3 we give details for the general case of maintaining matroid intersections. Finally, we close with some remarks and future directions in §4.

1.2 Related work

To our knowledge, online matroid intersection maintenance with recourse has not been studied previously. The special case of online bipartite matching problem with recourse was defined by [GKKV95], who gave an $\Omega(n \log n)$ lower bound. [CDKL09] gave optimal algorithms with $O(n \log n)$ recourse when clients arrive in random order, or when the graph is a forest. For the case of forests, [BLSZ18, BLSZ22] studied the shortest augmenting path algorithm and showed it to also be optimal. The first breakthrough on the general case of matching maintenance was by [BLSZ14] who gave a $O(n^{1.5})$ recourse bound; eventually [BHR18] gave the current best $O(n \log^2 n)$ bound.

The problem of load-balancing with recourse is closely related: [AGZ99, PW98, Wes00] show how to allocate jobs to machines and maintain near-optimum load while reassigning $O(\log n)$ jobs per timestep. [ABK94, AKP$^+$93, ANR92] show results for dynamic settings without reassignments, and observe strong lower bounds. [GKS14] show how to allocate unit jobs to machines in a restricted machines setting to maintain a load of $(1+\varepsilon)L$ with $O(1/\varepsilon)$ recourse; they give results for a dynamic flow variant. Recently similar results were given by [KLS22] for the case of unrelated machines, with logarithmic recourse. Very recently [BBLS23] studied a more general setting in which covering-packing constraints arrive and depart online and should be satisfied upon arrival. This setting captures as a special case a fully dynamic fractional load balancing/matching problem in which jobs arrive and depart online. They obtained an $O(\log(n/\varepsilon)/\varepsilon)$-competitive algorithm when the algorithm is given a $(1+\varepsilon)$ resource augmentation.

Several works [BHK09, EFN23] have also modelled recourse in online matching with *buybacks* or *cancellations*. In these settings, there is instead a penalty for recourse; for every online vertex that is matched the algorithm earns money, but the algorithm may choose to "buy back" resources from offline vertices, incurring a penalty. The buyback setting has also been extended beyond matchings to matroid and matroid intersection constraints [AK09, BHK09, BV11].

There is an enormous body of work on online bipartite matching problems *without* recourse, starting with the seminal work of Karp, Vazirani and Vazirani [KVV90]. In these settings, the algorithm makes irrevocable decisions, and the goal is to maximize the size/weight of the matchings; see, e.g., [Meh13, EIV23]. An extension of this to matroid intersection was studied by [GS17], who considered two arbitrary matroids defined on the same ground set whose elements arrive one at a time in a random order, and must be irrevocably picked/discarded, to maximize the size of the independent set selected. Another large body of work studies the min-weight perfect matching problem (mostly in metric settings); see, e.g., [MNP06, BBGN14, Rag18, PS21]. The techniques in these works are orthogonal to ours.

Our combinatorial decomposition for matroids produces a *matroid intersection skeleton* extending that for matching; this decomposition for matching was studied by [BHR18], and previously, under the name of *matching skeletons* by [GKK12, LS17] with the goal of understanding streaming algorithms for matchings. The matching skeleton was also used to derive an optimal competitive ratio in the batch arrival model of online bipartite matching [FN20, FNS21]. To the best of our knowledge, the extension to matroids has not been studied before; making further connections to streaming algorithms for matroid intersection remains an interesting future direction.

As discussed above, our work makes a connection to and builds on basic results on market equilibria [AD54, EG59, BS05]. Market equilibria and especially the design of algorithms for computing these equilibria have been the subject of intense study by the algorithmic game theory community over the last two decades. For an introduction to the topic, see chapters 5 and 6 of [NRTV07]. Of particular relevance to us is the work of Jain and Vazirani [JV07] on Eisenberg-Gale markets [EG59].

Convex programming techniques, and in particular the Eisenberg-Gale "fair allocation" convex program have also been used to guide combinatorial algorithms before, e.g., in the context of flow-time scheduling [IKM18, GGKS19]. However, these prior works do not consider the cost of recourse; they use the convex program to directly schedule jobs. We instead use it to compute prices and show the existence of short augmenting paths.

Note that while the problem we study can viewed as a dynamic graph problem, the cost function we study (bounding recourse) is unrelated to the kinds of cost functions studied in the dynamic graph algorithms literature.

2 Maintaining Matchings via Markets

We now present our market equilibria-based perspective for the bipartite matching case; we build on this for general matroids in §3. For matchings, the adversary fixes a bipartite graph (B, T, E) with n buyers B and m items T. The vertices in T (the offline side) are known up-front, whereas the vertices in B (and the edges between them) are revealed online (we can assume that the maximum matching after i arrivals has size i, and hence the maximum matching has size $n = |B|$; this is without loss of generality, see [BHR18, Obs. 9]). We see the edges between the i^{th} buyer (also called i) and its neighbors $N(i)$ only at time i.

If M_{i-1} is the maximum matching maintained by the algorithm after seeing $i-1$ vertices, and buyer i arrives, the *shortest augmenting path* algorithm (SAP) finds an (arbitrary) shortest augmenting path from i to a free item (if such a path exists), and augments the matching M_{i-1} along this path to get M_i.

Let ℓ_i denote the length of this shortest augmenting path found by the algorithm, and the goal is to bound the worst-case value of $\sum_{i=1}^{n} \ell_i$. There exist instances for which $\sum_{i-1}^{n} \ell_i = \Omega(n \log n)$ [GKKV95, Thm. 1]; the following result of [BHR18]—which we prove using market equilibria in this section—matches this lower bound up to a logarithmic factor.

Theorem 2.1. *The* Shortest Augmenting Path *(SAP) algorithm performs $O(n \log^2 n)$ changes.*

2.1 Preliminaries on the Fisher Market In the *Fisher's linear model* [BS05], we have n buyers and m divisible goods. Each buyer i has a *budget* m_i. The utility functions are linear: buyer i derives a utility $u_{ij}y_{ij}$ out of being allocated an amount y_{ij} of good j. There exist "market-clearing" prices for the goods and a corresponding equilibrium allocation of the goods to buyers in which (i) each good with a positive price is fully sold; (ii) buyers are only allocated goods that maximize their utility-per-price, sometimes called "bang-per-buck" (if a good has price 0, no buyer has positive utility for it, so we assume that each item has an interested buyer); and (iii) each buyer spends their entire budget.

Eisenberg and Gale [EG59] showed how to compute the market equilibrium allocation and prices in the Fisher model using the following convex program.

$$\text{(EG1)} \qquad \begin{aligned} \max \quad & \textstyle\sum_{i=1}^{n} m_i \ln \left(\sum_{j=1}^{m} u_{ij} y_{ij} \right) \\ \text{s.t.,} \quad & \textstyle\sum_{i=1}^{n} y_{ij} \le 1 \qquad\qquad \forall j = 1, 2, \ldots, m \\ & y_{ij} \ge 0. \end{aligned}$$

Let $\{y_{ij}\}$ be the optimal solution to the convex program. The Lagrangian dual variable p_j for each item j in the program can be interpreted as the price of that item. We imagine that the budget of a buyer when it arrives is 1 (and 0 before it arrives) and the utilities are $u_{ij} = 1$ for all edges $(i, j) \in E$ and 0 otherwise. The following properties can be derived from the KKT optimality conditions (see, e.g., [NRTV07, Chapter 5]).

THEOREM 2.2. *Let $\{p_j\}_{j=1}^m$ be an optimal dual solution to (EG1). Then the following hold:*
(i) All items are fully allocated: $\sum_{i=1}^n y_{ij} = 1$ for all j.
(ii) Buyers are buying the cheapest price items: $y_{ij} > 0 \Rightarrow p_j = \min_{j' \in N(i)} p_{j'}$.
(iii) Each buyer i spends all of their money: $\sum_{j=1}^m p_j \cdot y_{ij} = 1$, and hence $\sum_{j=1}^m p_j = n$.

Thus, with unit utilities each buyer i buys only the lowest-price items from $N(i)$; we refer to this lowest price for buyer i as q_i. Let us prove an additional important property of market-clearing prices for this utility function.

LEMMA 2.1. *Let i be the i^{th} buyer to arrives and let p, p' be the equilibrium prices before and after it is added. Let $q := \min_{j \in N(i)}\{p_j\}$. Then*

$$
\begin{aligned}
p'_j &= p_j \quad \text{if } p_j < q \\
p'_j &\geq p_j \quad \text{if } p_j \geq q.
\end{aligned}
\tag{2.1}
$$

Proof. Let $B_{\geq q}$ (resp. $T_{\geq q}$) be the set of buyers that buy at price at least q (resp. the set of items whose price is at least q) in the market equilibrium immediately prior to the arrival of buyer i. No buyer in $B_{\geq q}$ has an edge to an item in $T_{<q}$, since it would buy such an item otherwise. So if we find a market equilibrium for the subproblem consisting of buyers in $B_{\geq q} \cup \{i\}$ and show that all prices do not decrease in that equilibrium, we will be done: that equilibrium together with the equilibrium for $B_{<q}$ satisfies KKT conditions and hence is the new equilibrium.

Hence, let us restrict our attention to buyers in $B_{\geq q} \cup \{i\}$ and their neighbors $T_{\geq q}$. Towards a contradiction, let $T_< := \{j \mid p'_j < p_j\}$ be the subset of items whose price decreases in the new market equilibrium. Let $B_< := \{\ell \mid y_{\ell j} > 0, j \in T_<\}$ be the buyers who buy items from $T_<$. Since each such buyer expends all their budget and all goods are completely sold, we have $\sum_{j \in T_<} p_j \leq |B_<|$. However, after the price update *all* the minimum price items for buyers in $B_<$ must be in $T_<$. Moreover, these item prices strictly decreased, and other item prices may increase or remain the same. Thus, it must be that $\sum_{j \in T_<} p_j > \sum_{j \in T_<} p'_j \geq |B_<|$, where the second inequality holds since the buyers now spend all their budgets on items from $T_<$. This is a contradiction. ∎

2.2 The Expansion Lemma A *tail augmenting path* is an alternating path that starts from an arbitrary item, alternates between matched and unmatched edges, and ends in a free item. The *length* of a (tail) augmenting path is the number of items on the path. The following key lemma relates the length of these paths to the item prices.

LEMMA 2.2. (EXPANSION LEMMA) *Let (B, T, E) be a bipartite graph, let $M^\star$ be an arbitrary matching, and let p be the market clearing prices. Then, for any item j^* with $p_{j^*} \in [0, 1)$, there is a tail augmenting path from j^* whose length is $O\big(\frac{\ln n}{1 - p_{j^*}}\big)$.*

Proof. Consider market clearing prices p and an allocation y_{ij}. Let $j^* \in T$ be an arbitrary item with price $p_{j^*} \in [0, 1)$. We will denote $M^\star(i)$ to be the item that buyer i is matched to under $M^\star$. If $p_{j^*} = 0$ then there is no i such that $j^* \in N(i)$, the item is therefore unmatched, and the claim holds. Otherwise, we define the following sets inductively,

$$
\begin{aligned}
R_1 &= \{j^*\} \\
L_k &= \{i \mid M^\star(i) \in R_k\} \qquad && k = 1, 2, \ldots, \\
R_{k+1} &= \{j \mid \ y_{ij} > 0, i \in L_k\} \qquad && k = 1, 2, \ldots,
\end{aligned}
$$

That is, L_k is the set of buyers that are matched in $M^\star$ to items in R_k, and R_{k+1} is the set of items that are bought by some buyer in L_k in the market clearing allocation. We prove the following inductively: If $R_1, \ldots, R_k$ do not contain an unmatched item, then (a) $p_j \leq p_{j^*}$ for all items $j \in R_{k+1}$, and (b) the size $|R_{k+1}| \geq (1/p_{j^*})^k$.

The base case for R_1 trivially holds. We begin by proving property (a) of the induction. Consider an item $j \in R_{k+1}$. Since it is bought strictly positively by some item $i \in L_k$, its price p_j must be the price paid by i, that is, $p_j = q_i$. Moreover, observe that $q_i \le p_{M^\star(i)}$, by definition. And since $i \in L_k$, we have $M^\star(i) \in R_k$, so the induction hypothesis implies $p_{M^\star(i)} \le p_{j^*}$. Together, this shows that $p_j = q_i \le p_{M^\star(i)} \le p_{j^*}$ as desired.

To prove property (b) of the induction, suppose that R_k has no unmatched items, then we have that $|L_k| = |R_k|$. Hence, we have,

$$|R_{k+1}| \triangleq |\{j \mid \; y_{ij} > 0, i \in L_k\}| = \textstyle\sum_{j\in R_{k+1}} \sum_{i:j\in N(i)} y_{ij} \tag{2.2}$$

$$\ge \textstyle\sum_{i\in L_k} \sum_{j\in R_{k+1}} y_{ij} \tag{2.3}$$

$$= \textstyle\sum_{i\in L_k} \frac{1}{q_i} \ge \frac{|L_k|}{p_{j^*}} = \frac{|R_k|}{p_{j^*}} \ge \left(\frac{1}{p_{j^*}}\right)^k. \tag{2.4}$$

Equality (2.2) holds since every item j that is allocated is fully sold, and so $\sum_{i:j\in N(i)} y_{ij} = 1$, and (2.3) holds since the RHS sums only on the subset of $y_{ij} > 0$ from L_k to R_{k+1}. Finally, (2.4) holds by Theorem 2.2 (ii),(iii) since each buyer $i \in L_k$ spends all its money on items of price q_i that are in R_{k+1} and so $q_i \cdot \sum_{j\in R_{k+1}} y_{ij} = 1$. The next inequality holds since $q_i \le p_{j^*}$, and finally we use the induction hypothesis.

By our construction, if R_k contains an unmatched item j, then there is a tail augmenting path from j^* to j of length k. To conclude the proof, note that for any R_k whose items are all matched, $|R_k| \le n$ (the number of buyers), since otherwise there must be an unmatched item in R_k. Thus, $(\frac{1}{p_{j^*}})^{k-1} \le |R_k| \le n$. Simplifying we get that the length of such a tail augmenting path is at most $O\big(\max\{1, \frac{\ln(n)}{\ln(1/p_{j^*})}\}\big) = O\big(\frac{\ln n}{1-p_{j^*}}\big)$, where the final inequality follows from $1 - x \le \ln {}^1\!/\!_x$ for $x \in (0, 1]$. ∎

2.3 Bounding the Augmentations

FACT 2.1. *The price of all items* before the arrival *of the i^{th} buyer is either* 1 *or at most* $1 - \frac{1}{i}$. *There is no tail augmenting path from items with price* 1.

Proof. Recall that we are assuming that all buyers can be matched, hence all prices being at most 1 follows from the KKT optimality conditions for (EG1) and Hall's theorem. Let T_p be a set of items with some price $p \le 1$, and let $B_p := \{i \mid y_{ij} > 0, j \in T_p\}$. Then, since the buyers in B_p buy only items in T_p, we have $|B_p| = \sum_{i\in B_p, j\in T_p} p \cdot y_{ij} = p \cdot |T_p|$, giving $p = {}^{|B_p|}\!/\!_{|T_p|}$. This fraction is either 1, or at most $\frac{i-1}{i}$ (because $|B_p| \le i - 1$ before the arrival of i). For the items of price 1, the buyers in B_1 have edges only to items in T_1 (since otherwise, they would buy cheaper items) and $|B_1| = |T_1|$. Thus, by Hall's theorem, they are all matched and there can be no tail augmenting path from such items. ∎

THEOREM 2.1. *The* Shortest Augmenting Path *(SAP) algorithm performs $O(n \log^2 n)$ changes.*

Proof. For each item j, let $p_j(i)$ be its price *before* the arrival of the i^{th} buyer. Let $j_{\min} \in N(i)$ be a neighbor of i of minimal price, and define $q_{\min}(i) := \min_{j\in N(i)}\{p_j(i)\}$. By Fact 2.1 and the assumption that each arriving buyer can be matched, $q_{\min}(i) \le 1 - 1/n$.

Let $\Delta p_j(i)$ be the change in the price of item j due to the arrival of buyer i. By Lemma 2.2 the length of the shortest augmenting path from i is at most that of the tail augmenting path from $j_{\min}$ which can be bounded as follows,

$$\ell_i = O\left(\frac{\ln n}{1 - q_{\min}(i)}\right) = O\left(\frac{\ln n}{1 - q_{\min}(i)}\right) \sum_{j:p_j(i)\ge q_{\min}(i)} \Delta p_j(i) \le O(\ln n) \cdot \sum_{j\in T} \frac{\Delta p_j(i)}{1 - p_j(i)}, \tag{2.5}$$

where the second equality uses Theorem 2.2 (iii) to infer that $\sum_{j=1}^m p_j$ equals the current number of matched buyers, and Lemma 2.1 says that the price of items with $p_j < q_{\min}(i)$ do not change, which together imply that $\sum_{j:p_j\ge q_{\min}(i)} \Delta p_j(i) = 1$.

Hence, the total number of augmentations is

$$\sum_{i=1}^n \ell_i \le O(\ln n) \cdot \underbrace{\sum_i \sum_{j\in T} \frac{\Delta p_j(i)}{1 - p_j(i)}}_{(\star)} \le O(\ln n) \cdot n \cdot \left(1 + \int_{p=0}^{1-1/n} \frac{dp}{1-p}\right) = O(n \ln^2 n),$$

where the second inequality uses two facts: firstly, because $\sum_{j \in T} \Delta p_j(i) = 1$ and as $1/1-x$ is monotonically increasing, the sum $(\star)$ is maximized when n items have a final value of 1. Secondly, we bound the last term of any one sum $\sum_i \frac{\Delta p_j(i)}{1-p_j(i)}$ by 1, and the previous terms—for which $p_j(i) \leq 1 - 1/n$—by the integral (note that $\Delta p_j(i) \geq 0$ for all i, j by Lemma 2.1). This completes the proof. □

3 Extending to a General Matroid

We now prove our main result, for the intersection of a partition matroid $\mathcal{P}$ with a general matroid $\mathcal{M}$ over the same ground set E. The matroid $\mathcal{P}$ has n parts, each with rank 1. The elements of the i^{th} part are revealed at time i (the order of the parts is unknown). We assume that there is an independent set of size n after the final arrival[2], and hence the maximum independent set after the arrival of i parts has size i.

The *shortest augmenting path* also works in this setting. Let I be the chosen independent set before the arrival of i, and $\mathcal{P}|_i$ and $\mathcal{M}|_i$ be the matroids restricted to elements $E|_i$, the elements which have been revealed thus far. Define the *exchange graph* $D_{\mathcal{P}|_i,\mathcal{M}|_i}(I)$ as the bipartite graph on nodes $(I, E|_i \setminus I)$ with (directed) arcs:

1. $y \to x$ is an arc of $D_{\mathcal{P}|_i,\mathcal{M}|_i}(I)$ if $I - y + x$ is independent in $\mathcal{P}$.
2. $y \leftarrow x$ is an arc of $D_{\mathcal{P}|_i,\mathcal{M}|_i}(I)$ if $I - y + x$ is independent in $\mathcal{M}$.

The algorithm finds a shortest path in $D_{\mathcal{P}|_i,\mathcal{M}|_i}(I)$ from some element in P_i to a *free element* in $\mathcal{M}$ (that is, some $e \notin I$ for which $I + e$ is independent in $\mathcal{M}$). This is an *augmenting path*: it defines a valid sequence of exchanges to form a new independent set, and the resulting independent set will have size $|I| + 1$. The correctness of this algorithm and its analysis in the offline setting are due to Aigner and Dowling, and also Lawler [AD71, Law75] (see also [Sch03, §41.2]).

Again, let ℓ_i be the length of the shortest augmenting path upon the arrival of the ith part, and we want to bound the worst-case value of $\sum_{i=1}^{n} \ell_i$. We restate our main Theorem:

THEOREM 1.1. (MAIN THEOREM) *The* Shortest Augmenting Path *(SAP) algorithm for the* Online Matroid Intersection Maintenance *problem results in at most $O(n \log^2 n)$ total reassignments, where n is the rank of the intersection $\mathcal{M} \cap \mathcal{P}$.*

The remainder of the paper is dedicated to the proof of Theorem 1.1. It proceeds analogously to our proof for bipartite matchings in §2: we begin by defining a corresponding market we call the *matroid intersection market* which, in conjunction with the matroid intersection skeleton, yields "prices" for the elements of E. We then use properties of these prices to prove an "Expansion Lemma" generalizing Lemma 2.2. This lemma bounds the length of augmenting paths in terms of prices, and and hence gives our main result. However, as mentioned in §1, each of these steps requires us to build on the ideas we used for the matchings case.

3.1 The matroid intersection (MI) market The *matroid intersection (MI) market* is defined for an arbitrary matroid $\mathcal{M} = (E, \mathcal{I})$ and a partition matroid $\mathcal{P}$ on the elements in E. There is a set B of *buyers*, where each part P_i, $i = 1, \ldots, |B|$ is associated with a buyer i. In this market, E is the set of *items*, the items in P_i are precisely those that are of interest to buyer i, and each buyer arrives at the market with a budget of m_i dollars. The utility u_i to agent i for allocation $\{y_e\}_{e \in P_i}$ is $\sum_{e \in P_i} y_e$, and the allocation constraints are $\sum_{e \in S} y_e \leq \text{rank}_{\mathcal{M}}(S)$ for each $S \subseteq E$. It is immediate that the Fisher market with $u_{ij} \in \{0, 1\}$ is the special case in which $\mathcal{M}$ is a partition matroid. A market equilibrium for an MI market can be computed using the following convex program that optimizes over the matroid polytope of $\mathcal{M}$:

$$
\text{(EG2)} \qquad \begin{array}{lll} \max & \sum_{i \in B} m_i \cdot \log\left(\sum_{e \in P_i} y_e\right) & \\ & \sum_{e \in S} y_e \leq \text{rank}_{\mathcal{M}}(S) \quad \forall S \subseteq E & (\alpha_S) \\ & y_e \geq 0. & \end{array}
$$

[2]This is w.l.o.g.: if an arriving part does not cause an augmentation, then no future augmenting path will pass through that part. And thus, the part can be ignored henceforth. See also [BHR18, Obs. 9] for a description in the matching setting.

Here y_e is the amount of item e allocated (to the buyer whose part contains e), and α_S for $S \subseteq E$ are a set of dual variables. The optimal dual is not necessarily unique. For any optimal dual solution α, we define *prices* as

$$p_e := \sum_{S \ni e} \alpha_S \quad \text{for each element } e \in E. \tag{3.6}$$

The KKT conditions for (EG2) are the following:

(A) Stationarity and complementary slackness:

$$p_e \geq \frac{1}{\sum_{e' \in P_i} y_{e'}} \quad \text{and} \quad y_e \cdot \left(p_e - \frac{1}{\sum_{e' \in P_i} y_{e'}} \right) = 0 \qquad \forall e \in E. \tag{3.7}$$

$$\alpha_S \cdot \left(\sum_{e' \in S} y_{e'} - \operatorname{rank}_{\mathcal{M}}(S) \right) = 0 \qquad \forall S \subseteq E. \tag{3.8}$$

(B) Primal feasibility and $\alpha_S \geq 0$ for all $S \subseteq E$ (dual feasibility).

LEMMA 3.1. *Let α_S be any optimal dual values for the convex program (EG2), and p_e the corresponding prices. The following hold:*

(a) All goods are maximally allocated: For every e in a part P_i with $m_i \neq 0$, we have $p_e > 0$, so there is some $S \ni e$ such that $\alpha_S > 0$, and the corresponding primal constraint is tight (i.e., $\sum_{e \in S} y_e = \operatorname{rank}(S)$).

(b) Each buyer i is only buying elements (i.e. $y_e > 0$) of minimum price $q_i := \min_{e \in P_i} p_e$ (or equivalently, highest bang-per-buck).

(c) Each buyer spends all of their money: $q_i \cdot \sum_{e \in P_i} y_e = \sum_{e \in P_i} p_e y_e = m_i$.

(d) $\sum_{S \subseteq E} \alpha_S \cdot \operatorname{rank}(S) = \sum_{i \in B} m_i$.

Proof. All parts follow directly from the above KKT optimality conditions. ∎

Jain and Vazirani [JV10] proposed a generalization of the Fisher linear market called *Eisenberg-Gale markets*, which capture many interesting markets, such as the resource allocation framework of Kelly [Kel97]. By definition, equilibria in these markets can be computed using an Eisenberg-Gale type convex program. One specific class of these markets are the so-called *submodular utility allocation (SUA)* markets, in which there are n buyers (where buyer i has budget m_i) and each buyer has an associated utility u_i. There are packing constraints on the utilities, which are encoded via a *polymatroid* function $\nu : 2^{[n]} \to \mathbb{R}_+$ (i.e., the function ν is submodular, monotone, and $\nu(\emptyset) = 0$). The corresponding SUA convex program is

$$\max \Big\{ \sum_i m_i \log u_i \mid \sum_{i \in S} u_i \leq \nu(S) \;\; \forall S \subseteq [n], u \geq 0 \Big\}. \tag{3.9}$$

Note that in the MI market, we have an allocation constraint for each subset of E, whereas in an SUA market there is an allocation constraint only for each subset $\cup_{i \in A} P_i$, where $A \subseteq [n]$. Nonetheless, an application of a continuous version of Rado's theorem [McD75] can be used to prove the following. (The proof is deferred to Appendix A.)

LEMMA 3.2. *The MI market is a submodular utility allocation market.*

3.2 The Matroid Intersection Skeleton

In this section, we give a combinatorial description of a specific set of prices of elements with respect to the buyers' budgets. We refer to this decomposition as the *matroid intersection (MI) skeleton* of the matroid intersection market. In Lemma 3.4 we show these prices correspond to a very specific optimal dual solution to the convex program (EG2). Hence, although optimal dual solutions are not unique in general, this allows us to focus on a unique set of duals, which we subsequently show can be related to a notion of *expansion*.

DEFINITION 3.1. Denote the budget for a subset of buyers $B' \subseteq B$ to be $m(B') := \sum_{i \in B'} m_i$. Also define the *neighborhood* of these buyers as all elements in all of their parts: $N(B') := \bigcup_{i \in B'} P_i$. With this, define *inverse expansion* of a set of buyers as

$$\operatorname{InvExp}_{\mathcal{M}}(B') := \frac{m(B')}{\operatorname{rank}_{\mathcal{M}}(N(B'))}$$

We now give an algorithm that outputs a nested family of elements $\varnothing = E_0 \subseteq E_1 \subseteq \ldots \subseteq E_L = E$, which we call the matroid intersection skeleton, as well as prices for each of the items in E. The algorithm, which is given as Algorithm 1, does the following: in each step ℓ, the algorithm finds the (inclusion-wise maximal) set $B_{\ell+1}$ of buyers having maximum inverse expansion ρ. The elements spanned by the neighborhood of those buyers are assigned a price of ρ. These elements are then contracted in $\mathcal{M}$, and the buyers $B_{\ell+1}$ are "peeled off", with the process then repeating on the contracted matroid with the remaining buyers. Defining E_ℓ to be the set of elements contracted until the ℓ^{th} step gives us the desired nested family $\varnothing = E_0 \subseteq E_1 \subseteq \ldots \subseteq E_L = E$. Moreover, the dual variables α_S can be chosen to be positive only on the sets in this nested family.

A similar nested family (of buyers, instead of elements) is introduced in an algorithm for finding equilibria in SUA markets [JV07]; however, such a decomposition is not uniquely extendable to all items in an MI market. Our MI skeleton is one such extension, which in particular, connects the market to a notion of expansion on the matroid intersection.

Henceforth, we will use $\mathcal{M}_\ell$ to denote the matroid contraction $\mathcal{M}/E_\ell$, and $\text{rank}^{(\ell)}$ and $\text{InvExp}^{(\ell)}$ to denote the rank and the inverse expansion with respect to this matroid $\mathcal{M}_\ell$.

Algorithm 1: The Matroid Intersection Skeleton

1.1 Initialize $E_0 \leftarrow \varnothing$, the contracted elements.
1.2 Initialize $B_0^{\text{rem}} \leftarrow B$ the remaining buyers.
1.3 **for** $\ell = 0, 1, \ldots$ *until* B_ℓ^{rem} *is empty* **do**
1.4 $\quad \rho \leftarrow \max_{B' \subseteq B_\ell^{\text{rem}}} \text{InvExp}^{(\ell)}(B')$.
1.5 $\quad$ Find the (unique) largest set of buyers $B_{\ell+1} \subseteq B_\ell^{\text{rem}}$ with inverse expansion $\text{InvExp}^{(\ell)}(B_{\ell+1}) = \rho$.
1.6 $\quad$ Consider the elements $S' := \text{span}_{\mathcal{M}_\ell}(N(B_{\ell+1}))$.
1.7 $\quad$ Set the prices of each $e \in S'$ to be $\widehat{p}_{\ell+1} := \rho$.
1.8 $\quad E_{\ell+1} \leftarrow E_\ell \cup S'$.
1.9 $\quad B_{\ell+1}^{\text{rem}} \leftarrow B_\ell^{\text{rem}} \setminus B_{\ell+1}$.

LEMMA 3.3. *The MI skeleton is well-defined, and $\widehat{p}_\ell > \widehat{p}_{\ell+1}$ for every ℓ.*

The proof of Lemma 3.3 uses the submodularity of the matroid rank function, and uncrossing arguments to show the uniqueness of the sets B_ℓ, and the fact that the densities are strictly increasing; we defer the formal argument to Appendix A. Instead we focus on showing that the prices defined by Algorithm 1 indeed give us optimal duals. The proof proceeds by induction on the total number of rounds L. We begin with the case $L = 1$ (where the "densest" set of buyers form the entire matroid).

CLAIM 3.1. (SINGLE ROUND) *Consider an instance in which the inverse expansion of the entire set of buyers, $\rho := \mathit{InvExp}_{\mathcal{M}}(B)$, is the maximum inverse expansion. Then any optimal primal-dual pair of solutions y^*, α^* for the convex program (EG2) satisfies $p_e^* := \sum_{S \ni e} \alpha_S^* = \rho$. That is, all prices with respect to the optimal duals α^* are equal to ρ. One such optimal dual solution sets $\widehat{\alpha}_E = \rho$ and $\widehat{\alpha}_S = 0$ for all other sets.*

Proof. Let $S^{\max} \subseteq E$ be the set of elements with price $p^*_{\max} := \max_{e \in E} p_e^*$. We claim that $S^{\max} = E$.

As a first step, we show that constraint corresponding to $S^{\max}$ in (EG2) is tight. This uses the observation that if the constraints for two sets are tight, then the constraints for their union and intersection are also tight due to submodularity of the matroid rank function. Now define:

$$S_e := \bigcap_{\substack{S \ni e \\ \alpha_S^* > 0}} S.$$

This set S_e is the smallest tight set containing $e \in S^{\max}$. Note that any $e' \in S_e$ has price at least that of e, and hence belongs to $S^{\max}$ by the maximality of e's price. In turn this implies that $S^{\max} = \cup_{e \in S^{\max}} S_e$; writing $S^{\max}$ as a union of tight sets shows its tightness.

Next, define $B^{\max}$ as the set of buyers buying from $S^{\max}$; we claim that $N(B^{\max})$ is also tight. Indeed, $N(B^{\max}) \subseteq S^{\max}$, since the buyers in $B^{\max}$ must buy at the least price in their neighborhood. All other items in $S^{\max} \setminus N(B^{\max})$ must have $y_e = 0$ since they are not sold. This means $y(N(B^{\max})) = y(S^{\max}) = r(S^{\max}) \geq r(N(B^{\max}))$; combining with the constraint in (EG2) means the last inequality is an equality.

Now, observe that for any subset K of buyers, we have

$$m(K) = \sum_{i \in K} q_i^* \sum_{e \in P_i} y_e \leq p_{\max}^* \sum_{i \in K} \sum_{e \in P_i} y_e \leq p_{\max}^* \operatorname{rank}(N(K)), \tag{3.10}$$

where the equality uses Lemma 3.1(c). In particular, this implies that $\operatorname{InvExp}_{\mathcal{M}}(K) \leq p_{\max}^*$ for every $K \subseteq B$, and so $\rho = \max_K \operatorname{InvExp}_{\mathcal{M}}(K) \leq p_{\max}^*$.

Now consider setting $K = B^{\max}$: the first inequality in eq. (3.10) holds at equality because each buyer in $B^{\max}$ buys at price $p_{\max}^*$, while the second holds at equality by the tightness of $N(B^{\max})$. So $\operatorname{InvExp}_{\mathcal{M}}(B^{\max}) = p_{\max}^*$, which shows that $\rho \geq p_{\max}^*$, and in fact equality holds.

Finally, for $K = B$, eq. (3.10) implies

$$m(B) \leq p_{\max}^* \operatorname{rank}(N(B)) = \rho \operatorname{rank}(N(B)) = m(B)$$

by the assumption that $\operatorname{InvExp}_{\mathcal{M}}(B) = \rho$. Therefore, setting $K = B$ gives equality throughout the sums in eq. (3.10). That is, $q_i^* = p_{\max}^*$ for all $i \in B$, which means that each buyer i buys at price $p_{\max}^*$. In particular, $B^{\max} = B$ and $S^{\max} = N(B) = E$.

The fact that $\widehat{\alpha}$ is an optimal dual follows from checking the KKT conditions with respect to y^*: Equation (3.7) is satisfied trivially, and Equation (3.8) is satisfied since we have shown above that E is a tight set. ∎

Claim 3.1 shows that the prices we computed (for the case of $L = 1$) indeed correspond to optimal duals. Next, we extend the proof to the case of general L, via an inductive argument.

Lemma 3.4. (Prices are Optimal) *Consider the prices $\widehat{p}_1, \ldots, \widehat{p}_L$ from Algorithm 1, and define*

$$\widehat{\alpha}_{E_L} := \widehat{p}_L$$
$$\widehat{\alpha}_{E_\ell} := \widehat{p}_\ell - \widehat{p}_{\ell+1} \qquad \ell = 1, \ldots, L-1$$

and $\widehat{\alpha}_S = 0$ for all other $S \subseteq E$; note that all $\widehat{\alpha}_S \geq 0$ by Lemma 3.3. Then the resulting $\widehat{\alpha}$ forms an optimal dual solution to (EG2); in particular, it satisfies the KKT optimality conditions.

Proof. We construct an optimal primal solution $\widehat{y}$, such that $(\widehat{y}, \widehat{\alpha})$ form a primal/dual pair satisfying the KKT conditions for the convex program in (EG2).

We proceed via induction on L, the number of steps in Algorithm 1. Claim 3.1 precisely proves the base case of $L = 1$. For our inductive hypothesis, assume the lemma holds true for Algorithm 1 having up to k iterations. We consider an instance that requires $k+1$ iterations of the Algorithm 1. Consider the two markets

- market $\mathbb{M}^{(1)}$ on buyers B_1, where all budgets m_i are equal to those from market $\mathbb{M}$, and with the matroid $\mathcal{M}$ restricted to the ground set $E^{(1)} := N(B_1)$. Let $y^{(1)}$ be an optimal primal solutions. By Claim 3.1, the vector $\widehat{\alpha}^{(1)}$ where $\widehat{\alpha}^{(1)}_{N(B_1)} = p_1$ (and all other sets have $\widehat{\alpha}$ value 0) is an optimal dual solution to market $\mathbb{M}^{(1)}$.
- market $\mathbb{M}^{(2)}$ on buyers $B \setminus B_1$, where all budgets m_i are equal to those from market $\mathbb{M}$, and with the matroid $\mathcal{M}/\operatorname{span}(E^{(1)})$ (contracting elements $\operatorname{span}(E^{(1)})$). We apply the inductive hypothesis on market $\mathbb{M}^{(2)}$: Let $y^{(2)}$ be an optimal primal solution and let $\widehat{\alpha}^{(2)}$ be the (optimal) dual solution arising from Algorithm 1.

We may consider a gluing of $y^{(1)}$ and $y^{(2)}$ to create a primal solution $\widehat{y}$

$$\widehat{y}_e := \begin{cases} y_e^{(1)} & \text{if } e \in E^{(1)} \\ 0 & \text{if } e \in \operatorname{span}(E^{(1)}) \setminus E^{(1)} \\ y_e^{(2)} & \text{otherwise} \end{cases}$$

The proof that $\widehat{y}$ and $\widehat{\alpha}$ satisfy the KKT conditions for market $\mathbb{M}$ is an application of the gluing lemma, Lemma A.1. ∎

Definition 3.2. (Canonical Duals/Prices) The duals and prices for (EG2) generated by Algorithm 1 are called *canonical duals and prices.*

Theorem 3.1 shows that canonical duals are optimal duals for the convex program (EG2). We can now extend Lemma 3.1 to record two more properties that are specific to canonical prices.

Theorem 3.1. (Two More Properties) *Let α_S be the canonical dual values for* (EG2) *with corresponding prices p_e, and let y be any optimal primal solution for* (EG2)*. The following hold:*

(e) Let $F \subseteq E$ be a maximum-price basis of $\mathcal{M}$. Then $\sum_{e\in F} p_e = m(B)$.

(f) For $B' \subseteq B$ and $q_{\max} := \max_{i\in B'} q_i$, consider the set of elements $S := \{e \in N(B') : y_e > 0\}$, and $E_>$ the elements with prices higher than $q_{\max}$. Then $\operatorname{rank}_{\mathcal{M}/E_>}(S) \geq \frac{m(B')}{q_{\max}}$.

Proof. We prove:

(e) Let F be a maximum-price basis (with respect to $\mathcal{M}$) of elements. We first claim that for any $S \subseteq E$, if $\alpha_S > 0$ (in particular, $S = E_\ell$ for some ℓ), then F contains $\operatorname{rank}(S)$ elements from S. Observe that for each E_ℓ from the MI skeleton, the greedy algorithm to compute F always considers all elements of E_ℓ before considering elements of $E \setminus E_\ell$ (those of lower price). Therefore, $|F \cap E_\ell| = \operatorname{rank}(E_\ell)$.

Finally, Lemma 3.1(d) and the above claim together imply the desired:

$$m(B) = \sum_{S:\alpha_S>0} \alpha_S \cdot \operatorname{rank}(S) = \sum_{\ell} \sum_{e\in F\cap E_\ell} \alpha_{E_\ell} = \sum_{e\in F}\sum_{S\ni e} \alpha_S = \sum_{e\in F} p_e.$$

(f) Denote $q_{\max} := \max_{i\in B'} q_i$. First observe that since $y_e > 0$ for $e \in S$ (lying in part P_i for some $i \in B'$), Lemma 3.1(b) implies that $p_e = q_i \leq q_{\max}$. Now consider Algorithm 1, in particular, the step where all elements with prices strictly greater than $q_{\max}$ have been contracted. This is exactly the matroid $\mathcal{M}/E_>$. At this point, the maximum inverse expansion is at most $q_{\max}$ (by Lemma 3.3). Hence, for any subset of buyers B' at this stage, we have

$$\frac{m(B')}{\operatorname{rank}_{\mathcal{M}/E_>}(M(B'))} \leq q_{\max}.$$

This is precisely the desired result.

□

Part (e) is used later in the proof of the main theorem to bound the total price, while part (f) gives a relation between the prices in an arbitrary set of buyers and its expansion (in some contraction of $\mathcal{M}$).

3.3 Monotonicity of Prices We can now show that the prices of items only increase as new buyers arrive. Showing this for the minimum prices that buyers observe follows from a property called *competition monotonicity* that holds for SUA markets (and hence for our MI market). However, for the proof of Theorem 1.1 we will need to use the fact that the canonical prices are monotone for *all* items (and not just for the price that each buyer buys at). Note that since the item prices are not unique, monotonicity of prices may not hold for all items.

For the purposes of this proof, we view the arrival of buyers as a change in their budgets. That is, the budget of all buyers that have arrived is 1 and the budget of all buyers that have not yet arrived is 0. Thus, the arrival of i^{inc} is represented by an increase in its budget from $m_{i^{\text{inc}}} = 0$ to $m_{i^{\text{inc}}} = 1$. Throughout this section, we assume that the MI skeleton immediately before i^{inc} arrives is given by the sequence of buyers $B_1, \ldots, B_L$, and the nested family of elements $E_1 \subsetneq \ldots \subsetneq E_L = E$, where price p_ℓ is assigned to the elements of $E_\ell \setminus E_{\ell-1}$, and each buyer $i \in B_\ell$ buys at price $q_i^{\text{old}} = p_\ell$ (i.e., p_ℓ is the minimum price element in P_i). We will use the notation $B_{\leq k} := \cup_{j=1}^k B_j$.

Theorem 3.2. (Monotonicity of Prices) *Let i^{inc} be an incoming buyer and let p^{old}, p^{new} be the canonical price vectors for the instances before and after its arrival. Then, for every item $e \in E$,*

$$p_e^{new} \geq p_e^{old} \tag{3.11}$$

Moreover, let $q := \min_{e\in P_{i^{inc}}}\{p_e^{old}\}$ be the minimum price that buyer i^{inc} observes before arriving. If $p_e^{old} < q$, then,

$$p_e^{new} = p_e^{old} \tag{3.12}$$

We will prove Theorem 3.2 in two steps: first, we show monotonicity of the prices buyers are buying at[3]. Second, we use the structure of the MI skeleton to extend this monotonicity to all elements.

LEMMA 3.5. *For any buyer i, let $q_i^{old} = \min_{e \in P_i} p_e$ be the price it buys at immediately before the arrival of i^{inc}, and similarly q_i^{new} the price it buys at immediately after. Then $q_i^{new} \geq q_i^{old}$.*

Proof. Recall that by Lemma 3.2, the matroid intersection market is an SUA market. Jain and Vazirani [JV07] proved the following monotonicity result about SUA markets:

FACT 3.1. (COMPETITION MONOTONICITY) *At the addition of a new buyer i^{inc}, no other buyer's utility u_i, defined in MI markets as $u_i := \sum_{e \in P_i} y_e$, increases.*

An immediate corollary of this fact and Lemma 3.1, part (c) is that for all buyers i that arrive before i^{inc}, we have $q_i^{\text{new}} \geq q_i^{\text{old}}$.

It remains to show that the minimum price of elements in $P_{i^{\text{inc}}}$ also does not decrease. To see this, suppose that (immediately before its arrival) i^{inc} is in B_k. Then, the MI skeleton implies that

$$\text{InvExp}^{(k)}(B_k) \geq \text{InvExp}^{(k)}(B_k \setminus i^{\text{inc}}), \tag{3.13}$$

but since $m(B_k) = m(B_k \setminus i^{\text{inc}})$, we must have $\text{rank}^{(k)}(N(B_k)) = \text{rank}^{(k)}(N(B_k \setminus i^{\text{inc}}))$, or i^{inc} would not be in B_k. Therefore, $P_{i^{\text{inc}}} \subseteq \text{span}(N(B_{\leq k} \setminus i^{\text{inc}}))$. Therefore, since the prices paid by buyers in $B_{\leq k} \setminus i^{\text{inc}}$ do not decrease after the arrival of i^{inc}, once all of these buyers are peeled off, i^{inc} must be as well, at a price no smaller than its price was before its arrival. ∎

Now we can prove monotonicity of the prices of all elements.

Proof. [Proof of Theorem 3.2] We first prove Equation (3.11) and specifically that $p_e^{\text{new}} \geq p_e^{\text{old}}$ for each $e \in E_\ell \setminus E_{\ell-1}$ by induction on ℓ. There are two cases: since e was contracted at step ℓ of Algorithm 1, we must either have that $e \in N(B_\ell)$, or $e \in \text{span}_{\mathcal{M}_{\ell-1}}(N(B_\ell)) \setminus N(B_\ell)$.

1. If $e \in N(B_\ell)$, then $e \in N(i)$ for some $i \in B_\ell$, and $p_e^{\text{old}} = p_\ell = q_i^{\text{old}}$. But then by Lemma 3.5,
$$p_e^{\text{new}} \geq q_i^{\text{new}} \geq q_i^{\text{old}} = p_e^{\text{old}}.$$

2. If $e \in \text{span}_{\mathcal{M}_{\ell-1}}(N(B_\ell)) \setminus N(B_\ell)$, then we have that $e \in \text{span}_{\mathcal{M}}(E_{\ell-1} \cup N(B_\ell))$, by the definition of matroid contraction. By induction and case 1 above, we know that all elements e' in $E_{\ell-1} \cup N(B_\ell)$ have new price at least $p_{e'}^{\text{new}} \geq p_{e'}^{\text{old}} \geq p_\ell = p_e^{\text{old}}$. Moreover, by design of Algorithm 1, we must have that after the arrival,
$$p_e^{\text{new}} \geq \min_{e' \in E_{\ell-1} \cup N(B_\ell)} p_{e'}^{\text{new}} \geq p_e^{\text{old}}$$
where the first inequality follows because once all elements in $E_{\ell-1} \cup N(B_\ell)$ have been assigned a price, then so have the elements in their span, including e.

It remains to prove Equation (3.12), that $p_e^{\text{new}} = p_e^{\text{old}}$ if $p_e^{\text{old}} < q$. Recall that when running Algorithm 1, immediately before i^{inc}'s arrival, i^{inc} is peeled off at step k with corresponding price p_k. We claim that this implies that in the MI skeleton with $m(i^{\text{inc}}) = 1$, all buyers of $B^{(1)} := B_{\leq k}$ are peeled off before any buyer $i \in B^{(2)} := B \setminus B^{(1)}$, and hence the prices of all elements with $p_e^{\text{old}} < p_k$ are unchanged.

To prove this, we consider two submarkets of the MI market after the arrival of i^{inc} : (1) market instance $\mathbb{M}^{(1)}$ with buyers $B^{(1)}$, elements $E^{(1)} := N(B^{(1)})$, with matroid $\mathcal{M}^{(1)} := \mathcal{M}|_{E^{(1)}}$ and partition matroid $\mathcal{P}|_{E^{(1)}}$, and (2) market instance $\mathbb{M}^{(2)}$ with buyers $B^{(2)} := B \setminus B^{(1)}$, $E^{(2)} := E \setminus \text{span}(E^{(1)})$, matroid $\mathcal{M}^{(2)} := \mathcal{M}/\text{span}(E^{(1)})$, and partition matroid $\mathcal{P}|_{E^{(2)}}$. Budgets in both markets are precisely those immediately after the arrival of i^{inc}. Note that deriving an MI skeleton from each of these markets separately yields optimal primal and dual solutions $y^{(1)}, \alpha^{(1)}$ (for market $\mathbb{M}^{(1)}$) and $y^{(2)}, \alpha^{(2)}$ (for market $\mathbb{M}^{(2)}$) where all prices in market $\mathbb{M}^{(1)}$ are at least p_k and all prices buyers in market $\mathbb{M}^{(2)}$ buy at are exactly as they were from the MI skeleton on that market immediately

[3]For a buyer i with a budget of 0, this price is simply the price of the cheapest element of P_i

before i^{inc} arrived. Note also that the elements in $\text{span}(E^{(1)}) \setminus E^{(1)}$ are not included in either market. To complete the proof then, it suffices to show that these two solutions can be "glued together" to obtain an optimal primal and dual solution to the full market after the arrival of i^{inc}, where the prices of all elements in $E^{(2)}$ are given by the optimal solution to market $\mathbb{M}^{(2)}$ and thus are unchanged relative to the prices before i^{inc} arrived. This is shown in Lemma A.1. ◻

3.4 The Expansion Lemma We now prove our expansion lemma, which relates the length of augmenting paths in the exchange graph to prices of elements.

LEMMA 3.6. *Let I^* be an arbitrary independent set in the intersection $\mathcal{P} \cap \mathcal{M}$, and let $p \in \mathbb{R}^{|E|}_{\geq 0}$ be the canonical prices. Then, for any element $e^* \in E$ with current price $p_{e^*} \in [0,1)$ such that $I^* \cup e^* \in \mathcal{P}$, there is an augmenting path from e^* of length at most $O\big(\frac{\ln n}{1-p_{e^*}}\big)$.*

Just as in the case of matchings, we define a sequence of sets of buyers $\{L_k\}_k$ and of items $\{R_k\}_k$ inductively, based on the independent set I^* and the prices p. Firstly, for an element $e \in E$, let $\text{circuit}(e, I^*) \subseteq I^*$ be the elements from I^* in the circuit formed by adding e to I^*. (If $I^* \cup \{e\} \in \mathcal{M}$ then $\text{circuit}(e, I^*) = \emptyset$.) Now, for any $S \subseteq E$, we define

$$\text{circuit}(S, I^*) := \bigcup_{e \in S} \text{circuit}(e, I^*).$$

Secondly, for a set $S \subseteq E$, let $S_{\leq \tau} := \{e \in S \mid p_e \leq \tau\}$ denote the set of elements in S which have price less than or equal to some threshold τ. With this, we may define our sequences of sets:

$$\begin{aligned} R_1 &:= \{e^*\} \\ L_k &:= \{i \in B \mid P_i \cap \text{circuit}(R_k, I^*)_{\leq p_{e^*}} \neq \emptyset\} \\ R_{k+1} &:= \{e \in N(L_k) \mid y_e > 0\}. \end{aligned}$$

In order to prove Lemma 3.6, we need the following expansion claim.

CLAIM 3.2. *If $R_1, \ldots, R_k$ do not contain a free element with respect to $\mathcal{M}$, then*
(a) $p_e \leq p_{e^*}$ for all $e \in R_{k+1}$, and
(b) $\text{rank}_{\mathcal{M}}(R_{k+1}) \geq (1/p_{e^*})^k$.

Proof. We proceed by induction on k. The base case is $k = 0$, in which case both properties are true for the singleton set $R_1 = \{e^*\}$. Now, to inductively prove property (a): consider an element $e \in R_{k+1}$, and say it belongs to part P_i. Since $y_e > 0$, Lemma 3.1(b) implies that $p_e = q_i$. Moreover, $i \in L_k$ by definition, so part P_i contains an element from $\text{circuit}(R_k, I^*)_{\leq p_{e^*}}$, which in particular implies that it contains an element with price at most p_{e^*} and so $q_i \leq p_{e^*}$. Putting these two facts together shows $p_e \leq p_{e^*}$, and hence part (a). To prove part (b): for brevity, define the sets

$$K := \text{circuit}(R_k, I^*) \qquad \text{and} \qquad K_{\leq} := \text{circuit}(R_k, I^*)_{\leq p_{e^*}}.$$

Since $K_{\leq} \subseteq K \subseteq I^*$, and I^* is independent in the matroid $\mathcal{I}$, we have $\text{rank}_{\mathcal{I}}(K_{\leq}) = |K_{\leq}|$. This means that each $i \in L_k$ has exactly one element in $P_i \cap K_{\leq}$, and therefore $|L_k| = |K_{\leq}|$. Let $E_{>}$ be all the elements of price greater than p_{e^*}, and consider the contraction $\mathcal{M}_{\leq} := \mathcal{M}/E_{>}$. We claim that

$$\text{rank}_{(\mathcal{M}_{\leq})}(R_k) \leq |K_{\leq}| = |L_k|. \tag{3.14}$$

Before we prove the claim in (3.14), let us use it to complete the inductive proof for part (b). Indeed, Theorem 3.1(f), along with the fact that no client of L_k buys at price greater than p_{e^*}, tells us that

$$\text{rank}_{(\mathcal{M}_{\leq})}(R_{k+1}) \geq \frac{|L_k|}{p_{e^*}}.$$

Using (3.14) now proves the second part of the inductive claim:

$$|R_{k+1}| \geq \text{rank}_{(\mathcal{M}_{\leq})}(R_{k+1}) \geq \frac{\text{rank}_{(\mathcal{M}_{\leq})}(R_k)}{p_{e^*}} \geq \left(\frac{1}{p_{e^*}}\right)^k.$$

Finally, it remains to prove the inequality in (3.14). For the sake of contradiction, suppose

$$|R_k| \geq \text{rank}_{(\mathcal{M}_{\leq})}(R_k) > |K_{\leq}| \geq \text{rank}_{(\mathcal{M}_{\leq})}(K_{\leq}).$$

By the matroid exchange property, there is some $e \in R_k$ for which $K_{\leq} \cup \{e\}$ is independent. But by the definition of matroid contraction, this implies that in our un-contracted matroid $\mathcal{M}$ we have

$$\text{rank}_{\mathcal{M}}(K_{\leq} \cup \{e\} \cup E_{>}) > \text{rank}_{\mathcal{M}}(K_{\leq} \cup E_{>}).$$

Since $K \subseteq K_{\leq} \cup E_{>}$, submodularity gives $\text{rank}_{\mathcal{M}}(K \cup e) > \text{rank}_{\mathcal{M}}(K)$. But this contradicts the fact that K is independent, while $K \cup e$ contains the circuit $\text{circuit}(e, I^*)$ by construction. ◻

Proof. [Proof of Lemma 3.6] Indeed, if R_k does contain a free element with respect to $\mathcal{M}$, then there is a path in the exchange graph from e^* to this free element of length at most $2k$. Indeed, each element $e \in R_k$ is contained in P_i for some buyer $i \in L_{k-1}$, and i has a neighbor $e' \in \text{circuit}(R_{k-1}, I^*)_{\leq p_{e^*}} \subseteq I^*$. This means that there is some $e'' \in R_{k-1}$ for which $I^* - e' + e'' \in \mathcal{M}$, and $I^* - e' + e \in \mathcal{P}$. Inductively, e'' is reachable from e^* by an alternating path of length $2(k-1)$, so therefore e is reachable by an alternating path of length $2k$. Such an alternating path to a free element implies that the *shortest* AP to a free element is of length no more than $2k$.

On the other hand, if R_k contains no free element, then we have $({}^1/{}_{p_{e^*}})^{k-1} \leq \text{rank}_{\mathcal{M}}(R_k) \leq |L_k| \leq n$. Taking logs, we get that k is at most $O(\max\{1, \frac{\ln n}{\ln(1/p_{e^*})}\}) = O(\frac{\ln n}{1-p_{e^*}})$ when the process ends with an augmenting path of length at most $2k$, hence proving Lemma 3.6. ◻

3.5 Bounding the Total Augmentation Cost

Proof. [Proof of Theorem 1.1] When the ith buyer arrives, say the minimum price neighbor before its arrival (i.e. while it has 0 budget) has price $q_{\min}$. Denote the maximum-price basis before arrival i as F_{i-1}, having elements of prices $f_1^{(i-1)} \geq \ldots \geq f_R^{(i-1)}$ (where R denotes the rank of $\mathcal{M}$), and similarly for F_i after arrival i. $\Delta f_z^{(i)} := f_z^{(i)} - f_z^{(i-1)} \geq 0$, by Theorem 3.2 (monotonicity of prices). Then, using the Expansion lemma (Lemma 3.6), the length ℓ_i of the shortest augmenting path upon arrival i is at most

$$\ell_i \leq O\left(\frac{\ln n}{1-q_{\min}}\right) = O\left(\frac{\ln n}{1-q_{\min}}\right) \cdot \sum_{z=1}^{R} \Delta f_z^{(i)} \leq O(\ln n) \cdot \sum_{z=1}^{R} \frac{\Delta f_z^{(i)}}{1-f_z^{(i-1)}}.$$

where the equality follows from Theorem 3.1(e) (since the total budget increases by 1 at the arrival of i), and the final inequality again from Theorem 3.2, since either $\Delta f_z^{(i)} = 0$, or $q_{\min} \leq f_z^{(i-1)}$.

Denote $\widehat{f}_z := \max\{f_z^{(i)} : f_z^{(i)} < 1\}$. Notice that by Theorem 3.1(e) at most n elements in F_n can have final price $f_z^{(n)} = 1$. The total cost of all augmenting paths is at most:

$$\begin{aligned}
\sum_{i=1}^{n} \ell_i &\leq O(\ln n) \sum_{z=1}^{R} \sum_{i=1}^{n} \frac{\Delta f_z^{(i)}}{1-f_z^{(i-1)}} \\
&\leq O(\ln n) \sum_{z=1}^{R} \left(\mathbb{1}_{(z \leq n)} + \int_{p=0}^{\widehat{f}_z} \frac{dp}{1-p}\right) \\
&= O(n \ln n) + O(\ln n) \sum_{z=1}^{R} \ln\left(\frac{1}{1-\widehat{f}_z}\right)
\end{aligned}$$

It remains to upper bound the right hand term by $O(n \ln^2 n)$. Here we give a more careful argument of this fact than in §2. We can compute the maximum value the sum could possibly attain, given certain constraints on the prices $\widehat{f}_z$. Observe that at all times, prices satisfy $p_e \leq \frac{n}{n+1}$ or $p_e = 1$ for every element e. Indeed, Algorithm 1 assigns e a price $p = \frac{m(B')}{\text{rank}^{(\ell)}(N(B'))}$ for some ℓ and some set B' of buyers. If $p < 1$, then $\text{rank}^{(\ell)}(N(B')) \geq m(B')+1$, so $p \leq \frac{m(B')}{m(B')+1} \leq \frac{n}{n+1}$. Also from part (e) of Theorem 3.1, $\sum_z \widehat{f}_z \leq \sum_z f_z^{(n)} = n$. Now, we have

$$\left\{\begin{array}{ll} \max_p & \sum_{z=1}^{R} \ln\left(\frac{1}{1-p_z}\right) \\ \text{s.t.} & \sum_{z=1}^{R} p_z \leq n \\ & p_z \leq \frac{n}{n+1} \end{array}\right\} = \left\{\begin{array}{ll} \max_p & \sum_{z=1}^{R} \ln\left(\frac{1}{1-p_z}\right) \\ \text{s.t.} & \sum_{z=1}^{R} p_z = n \\ & p_z \leq \frac{n}{n+1} \end{array}\right\} = (n+1)\ln(n+1).$$

The first equality follows from monotonicity of the objective. The second comes from the following mass-shifting argument: given two elements with prices p, q and $p \geq q$, let their sum be $p + q = D$. Then

$$\frac{d}{dp}\left[\log\left(\frac{1}{1-p}\right) + \log\left(\frac{1}{1-(D-p)}\right)\right] = \frac{1}{1-p} + \frac{1}{1-(D-p)} \geq 0.$$

In particular, we can increase the sum by moving price-mass from p_i to p_j for any $p_i < p_j$. Therefore the maximizing p has $p_z = \frac{n}{n+1}$ for $n+1$ elements, and $p_z = 0$ for the rest. Since our prices $\{\widehat{f}_z\}$ satisfy the constraints in the maximization problem above, we have shown $\sum_{i=1}^{n} \ell_i \leq n \ln n + (n+1)\ln(n+1)\ln n = O(n \ln^2 n)$. □

4 Closing Remarks

We gave the matroid intersection maintenance problem, where we maintain a common base in the intersection of a partition matroid $\mathcal{P}$ and another arbitrary matroid $\mathcal{M}$, where the parts of $\mathcal{P}$ appear online, such that the total number of changes performed is $O(n \log^2 n)$. This extends the previous result for the special case of the intersection of two partition matroids (i.e., bipartite matching). Our results were based on viewing the problem from the perspective of market equilibria, and using market-clearing prices to bound the lengths of augmenting paths. Several open problems remain: the most natural one is whether we can improve the bound to $O(n \log n)$, or give a better lower bound, even for the bipartite matchings case. Can a better bound be given for the fractional variant of the problem? Can the matroid constraints be generalized to broader sets of packing constraints (again in the fractional case)?

Finally, what can we say about intersections of two arbitrary matroids? This problem is hopeless in full generality because it captures the edge-arrival model in bipartite matchings, which has a simple $\Omega(n^2)$ lower bound. Are there other interesting special cases which avoid these lower bounds?

Acknowledgments Part of this work was conducted while A. Gupta and S. Sarkar were visiting the Simons Institute for the Theory of Computing; they thank the institute for its generous hospitality.

A Appendix

LEMMA 3.2. *The MI market is a submodular utility allocation market.*

Proof. In a sub-modular utility allocation (SUA) market, the utility of buyers are defined by sub-modular packing constraints. The convex program corresponding to this market is

$$\text{(A.1)} \qquad \begin{aligned} \max \quad & \sum_{i \in B} m(i) \cdot \log u_i \\ & \sum_{i \in B'} u_i \leq \nu(B') \quad \forall B' \subseteq B \qquad (\alpha_{B'}) \\ & u_i \geq 0 \end{aligned}$$

where ν is a sub-modular, monotone function. We consider the following SUA market. Define for a subset of buyers $B' \subseteq B$ the sub-modular monotone function

$$\nu(B') := \text{rank}(N(B')).$$

We will show that this SUA market and the matroid intersection market are equivalent, i.e., an optimal set of $y_e's$ can be transformed into an optimal solution in the SUA market, and vice versa. To see that a solution to the matroid intersection market can be turned into a solution to the SUA market with the same objective value, consider the following assignment of u_i's:

$$u_i := \sum_{e \in P_i} y_e.$$

Clearly, the u_i's are feasible, and the objective values are the same.

For the converse direction, we use the following theorem of McDiarmid, an extension of Rado's theorem

THEOREM A.1. (PROPOSITION 2C OF [McD75]) *Let $G = (X, Y, E)$ be a bipartite graph, and $\mathbb{P}$ a polymatroid on ground set Y associated to polymatroid function ρ. Let $u \in \mathbb{R}^+_X$, then u is linked onto some vector $y \in \mathbb{P}$ if and only if*

$$u(B') \leq \rho(N_G(B')) \qquad \forall B' \subseteq X. \tag{A.2}$$

In McDiarmid's language, any solution u in the SUA market (and therefore satisfying (A.2)) is "linked to" some y in the polymatroid defined by polymatroid function $\rho(A) = \text{rank}_{\mathcal{M}}(A)$ (so $\mathbb{P}$ is the matroid polytope of $\mathcal{M}$). This precisely gives us a solution of y_e's that lie in the matroid polytope of $\mathcal{M}$ such that $u_i = \sum_{e \in P_i} y_e$. ∎

LEMMA 3.3. *The MI skeleton is well-defined, and $\widehat{p}_\ell > \widehat{p}_{\ell+1}$ for every ℓ.*

Proof. In particular, we have the following:

(a) Each $B_{\ell+1}$ is uniquely defined.

(b) At each step ℓ, every remaining buyer $i \in B^{\text{rem}}_\ell$ has a nonempty neighborhood. That is, $P_i \setminus E_\ell \neq \varnothing$.

(c) The prices assigned in step ℓ are greater than those assigned in later steps. That is, $\widehat{p}_\ell > \widehat{p}_{\ell+1}$.

First, we state a fact that will be useful:

FACT A.1. *For non-negative a, b, c, d, p, if $\frac{a}{b} \leq p$ and $\frac{c}{d} \leq p$, then $\frac{a+c}{b+d} \leq p$. Moreover, if either of the first two inequalities are strict, then so is the third.*

(a) Let B', B'' be two sets of buyers at iteration ℓ with maximum inverse expansion ρ. We claim that $B' \cup B''$ has the same inverse expansion. Observe that

$$\begin{aligned} m(B' \cup B'') &\leq \text{rank}^{(\ell)}(N(B') \cup N(B'')) \cdot \rho \\ &\leq \left(\text{rank}^{(\ell)}(N(B')) + \text{rank}^{(\ell)}(N(B'')) - \text{rank}^{(\ell)}(N(B') \cap N(B''))\right) \cdot \rho \\ &\leq m(B') + m(B'') - m(B' \cap B'') \\ &= m(B' \cup B''), \end{aligned}$$

where the first inequality follows from maximality of ρ, the second by submodularity, and the third by our assumption on the inverse expansion of B' and B''. Therefore equality holds throughout, and $\text{InvExp}^{(\ell)}(B' \cup B'') = \rho$. So the union of all sets of maximum inverse expansion gives the unique largest set of maximum inverse expansion.

(b) Consider for contradiction the first iteration $\ell + 1$ at which some remaining buyer $i \in B^{\text{rem}}_{\ell+1}$ has empty neighborhood. That means $i \notin B_{\ell+1}$, while its neighbors $P_i \setminus E_\ell$ are in $\text{span}_{\mathcal{M}_\ell}(N(B_{\ell+1}))$. But then we could add i to $B_{\ell+1}$ without increasing the rank of its neighborhood. That is $\text{rank}^{(\ell)}(N(B_{\ell+1})) = \text{rank}^{(\ell)}(N(B_{\ell+1} \cup i))$, and thus $\text{InvExp}^{(\ell)}(B_{\ell+1} \cup i) \geq \text{InvExp}^{(\ell)}(B_{\ell+1}) = \rho$, a contradiction to the maximality of $B_{\ell+1}$.

(c) For contradiction, suppose that $\widehat{p}_{\ell+1} > \widehat{p}_\ell$. We will argue that B_ℓ is not the maximum inverse expansion set in $\mathcal{M}_{\ell-1}$. Since B_ℓ and $B_{\ell+1}$ are disjoint, we have

$$\text{InvExp}^{(\ell-1)}(B_\ell \cup B_{\ell+1}) = \frac{m(B_\ell) + m(B_{\ell+1})}{\text{rank}^{(\ell-1)}(N(B_\ell) \cup N(B_{\ell+1}))}.$$

Furthermore, recall that $\mathcal{M}_\ell = \mathcal{M}/E_\ell = \mathcal{M}_{\ell-1}/\text{span}_{\mathcal{M}_{\ell-1}}(N(B_\ell))$, so by the definition of matroid contraction,

$$\text{rank}^{(\ell)}(N(B_{\ell+1})) + \text{rank}^{(\ell-1)}(N(B_\ell)) = \text{rank}^{(\ell-1)}(N(B_\ell) \cup N(B_{\ell+1})).$$

We use this substitution to get our second equality:

$$\begin{aligned} \text{InvExp}^{(\ell-1)}(B_\ell \cup B_{\ell+1}) &= \frac{m(B_\ell) + m(B_{\ell+1})}{\text{rank}^{(\ell-1)}(N(B_\ell) \cup N(B_{\ell+1}))} \\ &= \frac{m(B_\ell) + m(B_{\ell+1})}{\text{rank}^{(\ell-1)}(N(B_\ell)) + \text{rank}^{(\ell)}(N(B_{\ell+1}))} \\ &> \widehat{p}_\ell. \end{aligned}$$

The last strict inequality follows from Fact A.1 and our assumption that $\widehat{p}_{\ell+1} > \widehat{p}_{\ell}$. This contradicts that $\widehat{p}_{\ell}$ is the maximum inverse expansion in $\mathcal{M}_{\ell-1}$.

∎

A.1 The Gluing Lemma Fix an MI market $\mathbb{M}$, and let $B_1, B_2, \dots, B_L$ be the sets of buyers peeled off in Algorithm 1. Fix some $1 \le r < L$, and define $B_{\le r} := \bigcup_{\ell \le r} B_\ell$, and $B_{>r} := B \setminus B_{\le r}$. Now we define two MI markets:

- market $\mathbb{M}^{(1)}$ on the ground set $E^{(1)} := N(B_{\le r})$, with matroid $\mathcal{M}^{(1)} := \mathcal{M}|_{E^{(1)}}$, and partition matroid $\mathcal{P}|_{E^{(1)}}$. The buyers are $B_{\le r}$, and their budgets are the same as in market $\mathbb{M}$. Let $y^{(1)}, \widehat{\alpha}^{(1)}$ be corresponding optimal primal and dual solutions, where $\widehat{\alpha}^{(1)}$ is the dual solution from the MI skeleton on market $\mathbb{M}^{(1)}$.
- market $\mathbb{M}^{(2)}$ on the ground set $E^{(2)} := E \setminus \mathrm{span}(E^{(1)})$, with matroid $\mathcal{M}^{(2)} := \mathcal{M}/\,\mathrm{span}(E^{(1)})$, and partition matroid $\mathcal{P}|_{E^{(2)}}$. The buyers are $B_{>r}$, and their budgets are the same as in market $\mathbb{M}$. Let $y^{(2)}, \widehat{\alpha}^{(2)}$ be corresponding optimal primal and dual solutions, where $\widehat{\alpha}^{(2)}$ is the dual solution from the MI skeleton on market $\mathbb{M}^{(2)}$.

Additionally, let $\widehat{p}_{r+1}$ be the price assigned in step $r+1$ of Algorithm 1 for market $\mathbb{M}$.

Lemma A.1. (Gluing Lemma) *Let y^+ be*

$$y^+_e := \begin{cases} y^{(1)}_e & \text{if } e \in E^{(1)} \\ 0 & \text{if } e \in \mathrm{span}(E^{(1)}) \setminus E^{(1)} \\ y^{(2)}_e & \text{otherwise} \end{cases}$$

and α^+ be

$$\alpha^+_S := \begin{cases} \widehat{\alpha}^{(1)}_{S'} & S = \mathrm{span}(S') \text{ for } S' \subsetneq E^{(1)} \text{ with } \widehat{\alpha}^{(1)}_{S'} > 0 \\ \widehat{\alpha}^{(1)}_{E^{(1)}} - \widehat{p}_{r+1} & S = \mathrm{span}(E^{(1)}) \\ \widehat{\alpha}^{(2)}_{S'} & S = S' \cup \mathrm{span}(E^{(1)}) \text{ for } S' \subseteq E^{(2)} \text{ with } \widehat{\alpha}^{(2)}_{S'} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Then, y^+ is an optimal solution to the convex program (EG2) for market instance $\mathbb{M}^+$, identical to market $\mathbb{M}$. And, in particular, the prices defined by α^+ are identical to those assigned in markets $\mathbb{M}^{(1)}$ and $\mathbb{M}^{(2)}$.

Moreover, the lemma still holds if the budgets of buyers in $B_{\le r}$ are allowed to increase from their original values in market $\mathbb{M}$, with the increase happening simultaneously in $\mathbb{M}^{(1)}$ and $\mathbb{M}^+$.

Proof. It suffices to check that y^+ and α^+ satisfy the KKT conditions for market $\mathbb{M}^+$.

Primal Feasibility: We first verify primal feasibility: for any $S \subseteq E$, we have

$$\sum_{e \in S} y^+_e = \sum_{e \in S \cap E^{(1)}} y^{(1)}_e + \sum_{e \in S \cap E^{(2)}} y^{(2)}_e \le \mathrm{rank}_{\mathcal{M}^{(1)}}(S \cap E^{(1)}) + \mathrm{rank}_{\mathcal{M}^{(2)}}(S \cap E^{(2)}) \le \mathrm{rank}(S)$$

where the first inequality follows from feasibility of $y^{(1)}$ and $y^{(2)}$ in their respective markets, and the second from the definition of matroid contraction.

Dual Feasibility: For all $S \ne \mathrm{span}(E^{(1)})$, it is clear that $\alpha^+_S \ge 0$, so dual feasibility holds. It remains to show for $S = \mathrm{span}(E^{(1)})$. We first observe that $\sum_{S' \subseteq E^{(2)}} \alpha^{(2)}_{S'} = \widehat{p}_{r+1}$, since Algorithm 1 on market $\mathbb{M}^{(2)}$ is identical to steps $r+1, \dots, L$ in market $\mathbb{M}$ (by construction). Moreover, $\widehat{\alpha}^{(1)}_{E^{(1)}}$ is the minimum price of any element in market $\mathbb{M}^{(1)}$. In the case that budgets are not increased, $\widehat{\alpha}^{(1)}_{E^{(1)}}$ is clearly equal to $\widehat{p}_r$, since the steps Algorithm 1 takes for market $\mathbb{M}^{(1)}$ are the same as the first r steps on market $\mathbb{M}$.

If budgets *have* been increased, then still we have $\widehat{\alpha}^{(1)}_{E^{(1)}} \ge \widehat{p}_r$, by the monotonicity of prices given in Lemma 3.5. In either case,

$$\widehat{\alpha}^{(1)}_{E^{(1)}} \ge \widehat{p}_r > \widehat{p}_{r+1}.$$

So $\alpha^+_{\mathrm{span}(E^{(1)})} = \widehat{\alpha}^{(1)}_{E^{(1)}} - \widehat{p}_{r+1} \ge 0$, and dual feasibility is satisfied.

KKT Stationarity Condition (eq. (3.7)): Let use define

$$p_e^+ := \sum_{S \ni e} \alpha_S^+.$$

and likewise $\widehat{p}_e^{(1)} = \sum_{S \ni e} \widehat{\alpha}_S^{(1)}$ for $e \in E^{(1)}$ and $\widehat{p}_e^{(2)} = \sum_{S \ni e} \widehat{\alpha}_S^{(2)}$ for $e \in E^{(2)}$. We first observe that $p_e^+ = \widehat{p}_e^{(1)}$ for every $e \in E^{(1)}$, and likewise for $e \in E^{(2)}$. This is immediate for every $e \in E^{(2)}$, by the definition of α^+. Moreover, since $\sum_{S' \subseteq E^{(2)}} \alpha_{S'}^{(2)} = \widehat{p}_{r+1}$, we have that for $e \in E^{(1)}$

$$p_e^+ = \sum_{S \ni e} \alpha_S^+ = p_{r+1} + \sum_{\substack{S \ni e \\ S \subseteq \mathrm{span}(E^{(1)})}} \alpha_S^+ = \sum_{\substack{S' \ni e \\ S' \subseteq E^{(1)}}} \widehat{\alpha}_{S'}^{(1)} = \widehat{p}_e^{(1)}$$

This implies that α^+ and y^+ satisfy the stationarity condition for $e \in E^{(1)} \cup E^{(2)}$, since $y^{(1)}$, $\widehat{\alpha}^{(1)}$, $y^{(2)}$, $\widehat{\alpha}^{(2)}$ are optimal and satisfy the stationarity conditions.

Lastly, we must check whether edges $e \in \mathrm{span}(E^{(1)}) \setminus E^{(1)}$ satisfy the stationarity condition. Say edge e belongs to buyer i's part. By definition, buyer i is a part of market $\mathbb{M}^{(2)}$. By the the optimality of $y^{(2)}$ and $\widehat{\alpha}^{(2)}$ in market $\mathbb{M}^{(2)}$, we know the price buyer i is purchasing at is precisely equal to

$$q_i = \frac{1}{\sum_{e' \in P_i} y_{e'}^+}.$$

Since buyer i is a part of market $\mathbb{M}^{(2)}$, we have $q_i \leq \widehat{p}_{r+1}$. By monotonicity of prices, all prices in market $\mathbb{M}^{(1)}$ are at least $\widehat{p}_{r+1}$. Notice that, by construction, p_e^+ has the same price as some element in $E^{(1)}$. Hence,

$$p_e^+ > \widehat{p}_{r+1} \geq q_i = \frac{1}{\sum_{e' \in P_i} y_{e'}^+}.$$

Note that the second part of (eq. (3.7)) is trivially satisfied since $y_e^+ = 0$.

KKT Complementary Slackness (eq. (3.8)): We show that sets with non-zero α_S^+ value are tight under y_e^+. We go through the cases on what values α_S^+ could take. If $\alpha_S^+ = \widehat{\alpha}_{S'}^{(1)}$ for some $S = \mathrm{span}(S')$ where $\widehat{\alpha}_{S'}^{(1)} > 0$ (case 1), then,

$$\sum_{e \in S'} y_e^{(1)} = \mathrm{rank}(S')$$

by complementary slackness on market $\mathbb{M}^{(1)}$. Therefore,

$$\begin{aligned} \sum_{e \in S} y_e^+ &= \sum_{e \in S} y_e^{(1)} + \sum_{e \in S \setminus S'} 0 \\ &= \mathrm{rank}(S') = \mathrm{rank}(S). \end{aligned}$$

Similar logic holds for the case $S = \mathrm{span}(E^{(1)})$ (case 2). Lastly, if $\alpha_S^+ = \widehat{\alpha}_{S'}^{(2)}$ for some $S = S' \cup \mathrm{span}(E^{(1)})$, then, since S' is disjoint from $\mathrm{span}(E^{(1)})$ and $\mathrm{span}(E^{(1)})$ is a tight set under y_e^+, we have

$$\begin{aligned} \sum_{e \in S} y_e^+ &= \sum_{e \in S'} y_e^+ + \sum_{e \in \mathrm{span}(E^{(1)})} y_e^+ \\ &= \sum_{e \in S'} y_e^{(2)} + \sum_{e \in \mathrm{span}(E^{(1)})} y_e^+ \\ &= \mathrm{rank}(S') + \mathrm{rank}(E^{(1)}) \\ &\geq \mathrm{rank}(S' \cup \mathrm{span}(E^{(1)})) \end{aligned}$$

which means $\sum_{e \in S} y_e^+ = \mathrm{rank}(S' \cup \mathrm{span}(E^{(1)}))$ as desired. ∎

References

[ABK94] Yossi Azar, Andrei Z. Broder, and Anna R. Karlin. On-line load balancing. *Theoretical Computer Science*, 130(1):73–84, 1994.

[AD54] Kenneth J. Arrow and Gérard Debreu. Existence of an equilibrium for competitive economy. *Econometrica*, pages 265–290, 1954.

[AD71] Martin Aigner and Thomas A. Dowling. Matching theory for combinatorial geometries. *Trans. Amer. Math. Soc.*, 158:231–245, 1971.

[AGZ99] M. Andrews, M. X. Goemans, and L. Zhang. Improved bounds for on-line load balancing. *Algorithmica*, 23(4):278–301, 1999.

[AK09] B. V. Ashwinkumar and Robert Kleinberg. Randomized online algorithms for the buyback problem. In Stefano Leonardi, editor, *Internet and Network Economics*, pages 529–536, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.

[AKP+93] Yossi Azar, Bala Kalyanasundaram, Serge A. Plotkin, Kirk Pruhs, and Orli Waarts. Online load balancing of temporary tasks. In *Proceedings of the 1993 Workshop on Algorithms and Data Structures*, 1993.

[ANR92] Yossi Azar, Joseph Naor, and Raphael Rom. The competitiveness of on-line assignments. In *SODA: ACM-SIAM Symposium on Discrete Algorithms (A Conference on Theoretical and Experimental Analysis of Discrete Algorithms)*, 1992.

[BBGN14] Nikhil Bansal, Niv Buchbinder, Anupam Gupta, and Joseph Naor. A randomized o(log2 k)-competitive algorithm for metric bipartite matching. *Algorithmica*, 68(2):390–403, 2014.

[BBLS23] Sayan Bhattacharya, Niv Buchbinder, Roie Levin, and Thatchaphol Saranurak. Chasing positive bodies. *CoRR, to appear in FOCS 2023*, abs/2304.01889, 2023.

[BHK09] Moshe Babaioff, Jason Hartline, and Robert Kleinberg. Selling ad campaigns: Online algorithms with cancellations. In *Proceedings 10th ACM Conference on Electronic Commerce (EC-2009), Stanford, California, USA, July 6–10, 2009*, pages 61–70, 07 2009.

[BHR18] Aaron Bernstein, Jacob Holm, and Eva Rotenberg. Online bipartite matching with amortized replacements. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 947–959. SIAM, 2018.

[BLSZ14] Bartlomiej Bosek, Dariusz Leniowski, Piotr Sankowski, and Anna Zych. Online bipartite matching in offline time. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 384–393. IEEE Computer Society, 2014.

[BLSZ18] Bartlomiej Bosek, Dariusz Leniowski, Piotr Sankowski, and Anna Zych-Pawlewicz. Shortest augmenting paths for online matchings on trees. *Theory Comput. Syst.*, 62(2):337–348, 2018.

[BLSZ22] Bartlomiej Bosek, Dariusz Leniowski, Piotr Sankowski, and Anna Zych-Pawlewicz. A tight bound for shortest augmenting paths on trees. *Theor. Comput. Sci.*, 901:45–61, 2022.

[BS05] William C. Brainard and Herbert E. Scarf. How to compute equilibrium prices in 1891. *The American Journal of Economics and Sociology*, 64(1):57–83, 2005.

[BV11] Ashwinkumar Badanidiyuru Varadaraja. Buyback problem - approximate matroid intersection with cancellation costs. In Luca Aceto, Monika Henzinger, and Jiří Sgall, editors, *Automata, Languages and Programming*, pages 379–390, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[CDKL09] Kamalika Chaudhuri, Constantinos Daskalakis, Robert D. Kleinberg, and Henry Lin. Online bipartite perfect matching with augmentations. In *INFOCOM 2009. 28th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 19-25 April 2009, Rio de Janeiro, Brazil*, pages 1044–1052. IEEE, 2009.

[CS12] Mosharaf Chowdhury and Ion Stoica. Coflow: a networking abstraction for cluster applications. In Srikanth Kandula, Jitendra Padhye, Emin Gün Sirer, and Ramesh Govindan, editors, *11th ACM Workshop on Hot Topics in Networks, HotNets-XI, Redmond, WA, USA - October 29 - 30, 2012*, pages 31–36. ACM, 2012.

[Edm65] Jack Edmonds. Minimum partition of a matroid into independent subsets. *J. Res. Nat. Bur. Standards Sect. B*, 69B:67–72, 1965.

[EFN23] Farbod Ekbatani, Yiding Feng, and Rad Niazadeh. Online resource allocation with buyback: Optimal algorithms via primal-dual. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, page 583, New York, NY, USA, 2023. Association for Computing Machinery.

[EG59] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.

[EIV23] Federico Echenique, Nicole Immorlica, and Vijay Vazirani, editors. *Online and Matching Based Market Design.* Cambridge University Press, 2023.

[FN20] Yiding Feng and Rad Niazadeh. Batching and optimal multi-stage bipartite allocations. *SSRN Electronic Journal*, 01 2020.

[FNS21] Yiding Feng, Rad Niazadeh, and Amin Saberi. Two-stage stochastic matching with application to ride hailing. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2862–2877. SIAM, 2021.

[GGKS19] Naveen Garg, Anupam Gupta, Amit Kumar, and Sahil Singla. Non-clairvoyant precedence constrained scheduling. In *ICALP*, pages 63:1–63:14, Jul 2019.

[GKK12] Ashish Goel, Michael Kapralov, and Sanjeev Khanna. On the communication and streaming complexity of maximum bipartite matching. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 468–485. SIAM, 2012.

[GKKV95] Edward F. Grove, Ming-Yang Kao, P. Krishnan, and Jeffrey Scott Vitter. Online perfect matching and mobile computing. In Selim G. Akl, Frank K. H. A. Dehne, Jörg-Rüdiger Sack, and Nicola Santoro, editors, *Algorithms and Data Structures, 4th International Workshop, WADS '95, Kingston, Ontario, Canada, August 16-18, 1995, Proceedings*, volume 955 of *Lecture Notes in Computer Science*, pages 194–205. Springer, 1995.

[GKS14] Anupam Gupta, Amit Kumar, and Cliff Stein. Maintaining assignments online: Matching, scheduling, and flows. In *SODA*, pages 468–479, 2014.

[GS17] Guru Prashanth Guruganesh and Sahil Singla. Online matroid intersection: Beating half for random arrival. In Friedrich Eisenbrand and Jochen Könemann, editors, *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings*, volume 10328 of *Lecture Notes in Computer Science*, pages 241–253. Springer, 2017.

[IKM18] Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Competitive algorithms from competitive equilibria: Non-clairvoyant scheduling under polyhedral constraints. *J. ACM*, 65(1):3:1–3:33, 2018.

[IMPP19] Sungjin Im, Benjamin Moseley, Kirk Pruhs, and Manish Purohit. Matroid coflow scheduling. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPIcs*, pages 145:1–145:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[JKR17] Hamidreza Jahanjou, Erez Kantor, and Rajmohan Rajaraman. Asymptotically optimal approximation algorithms for coflow scheduling. In Christian Scheideler and Mohammad Taghi Hajiaghayi, editors, *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2017, Washington DC, USA, July 24-26, 2017*, pages 45–54. ACM, 2017.

[JV07] Kamal Jain and Vijay V. Vazirani. Eisenberg-gale markets: algorithms and structural properties. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 364–373, 2007.

[JV10] Kamal Jain and Vijay V. Vazirani. Eisenberg-gale markets: Algorithms and game-theoretic properties. *Games Econ. Behav.*, 70(1):84–106, 2010.

[Kel97] Frank Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8(1):33–37, 1997.

[KLS22] Ravishankar Krishnaswamy, Shi Li, and Varun Suriyanarayana. Online unrelated-machine load balancing and generalized flow with recourse. *CoRR*, abs/2211.16216, 2022. STOC 2023, to appear.

[KVV90] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 352–358. ACM, 1990.

[Law75] Eugene L. Lawler. Matroid intersection algorithms. *Math. Programming*, 9(1):31–56, 1975.

[LS17] Euiwoong Lee and Sahil Singla. Maximum matching in the online batch-arrival model. In *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings*, volume 10328 of *Lecture Notes in Computer Science*, pages 355–367. Springer, 2017.

[McD75] Colin J. H. McDiarmid. Rado's theorem for polymatroids. *Math. Proc. Cambridge Philos. Soc.*, 78(2):263–281, 1975.

[Meh13] Aranyak Mehta. Online matching and ad allocation. *Found. Trends Theor. Comput. Sci.*, 8(4):265–368, 2013.

[MNP06] Adam Meyerson, Akash Nanavati, and Laura J. Poplawski. Randomized online algorithms for minimum metric bipartite matching. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 954–959. ACM Press, 2006.

[NRTV07] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[PS21] Enoch Peserico and Michele Scquizzato. Matching on the line admits no o($\sqrt{\text{log n}}$)-competitive algorithm. In *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 103:1–103:3, 2021.

[PW98] S. Phillips and J. Westbrook. On-line load balancing and network flow. *Algorithmica*, 21(3):245–261, 1998.

[Rag18] Sharath Raghvendra. Optimal analysis of an online algorithm for the bipartite matching problem on a line. In *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*, volume 99 of *LIPIcs*, pages 67:1–67:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[Sch03] Alexander Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. B*, volume 24 of *Algorithms and*

Combinatorics. Springer-Verlag, Berlin, 2003. Matroids, trees, stable sets, Chapters 39–69.
[Wes00] Jeffery Westbrook. Load balancing for response time. *J. Algorithms*, 35(1):1–16, 2000.