Online Matroid Embeddings

Andrés Cristi
EPFL, Lausanne, Switzerland

Paul Dütting
Google Research, Zürich, Switzerland

Robert Kleinberg
Cornell University, Ithaca, NY, USA
Google Research, New York, NY, USA

Renato Paes Leme
Google Research, New York, NY, USA

Neel Patel
Google Research, Zürich, Switzerland

Abstract

We introduce the notion of an online matroid embedding, which is an algorithm for mapping an unknown matroid that is revealed in an online fashion to a larger-but-known matroid. The existence of such embedding enables a reduction from the version of the matroid secretary problem where the matroid is unknown to the version where the matroid is known in advance. We establish the existence of such an embedding for binary matroids, and use it to relate variants of the binary matroid secretary problem to each other, showing that seemingly simpler problems are in fact equivalent to seemingly harder ones (up to constant-factors). Specifically, we show this to be the case for the version of the matroid secretary problem in which the binary matroid is not known in advance, and where it is known in advance. We also show that the version with known matroid structure is equivalent to the problem where weights are not fully adversarial but drawn from a known pairwise-independent distribution.

2012 ACM Subject Classification Theory of computation → Online algorithms; Theory of computation → Random projections and metric embeddings

Keywords and phrases Matroids, Secretary Problem, Online Algorithm

Digital Object Identifier 10.4230/LIPIcs.ICALP.2026.69

Category Track A: Algorithms, Complexity and Games

Related Version *Extended Version*: https://arxiv.org/abs/2407.10316

Funding *Andrés Cristi*: Partially funded by the Center for Mathematical Modeling (CMM), FB210005, BASAL funds for centers of excellence (ANID-Chile).

1 Introduction

A common setup in online algorithms is to have a matroid whose structure is revealed to the algorithm one element at a time. The algorithm processes the elements of the ground set in sequence, and at each point in time, it has access to the dependencies between the elements that have already arrived. Typical examples include the famous matroid secretary problem (MSP) [9, 7, 33, 22] and matroid prophet inequalities [15, 31].

In such problems, is there any advantage in knowing the matroid structure in advance? Imagine the following situation: we are processing an unknown matroid **M**; however, we know a fixed (potentially very large) matroid **BigM** that has an isomorphic copy of every possible matroid **M**, and we can construct this embedding online. We will show that if such an object exists, then we can reduce the version of the problem where the matroid is revealed

53rd International Colloquium on Automata, Languages, and Programming (ICALP 2026).
Editors: Sayan Bhattacharya, Danupon Nanongkai, Michael Benedikt, and Gabriele Puppis;
Article No. 69; pp. 69:1–69:24
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

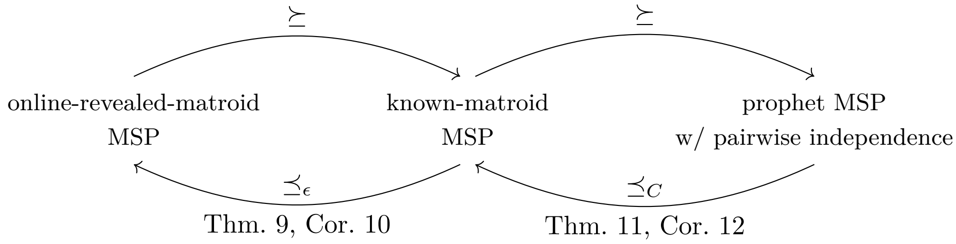

Figure 1 Reductions for binary matroids. We use $P \succeq Q$ to indicate that P is harder than Q, and we use $\succeq_\epsilon$ and $\succeq_C$ to designate an additive ε or multiplicative factor C loss in approximation.

online to the version of the problem where the matroid structure is known, by assuming our matroid is **BigM**. Moreover, if the on-the-fly embedding maintains uniform random order, then the existence of such an online embedding implies that in that class, the matroid secretary problem with *unknown* structure is no harder than the matroid secretary problem with *known* structure.

Important recent progress on the MSP has established that it is equivalent to the matroid prophet secretary problem with correlated distributions [18, 17]. Though seemingly unrelated, another consequence of the existence of such online embeddings will be that – for certain matroids – this equivalence holds even if we impose pairwise independence.

1.1 Our Contribution

Online Matroid Embedding. Our main conceptual contribution is to define the notion of an *online matroid embedding* (OME) in Section 3. For a given class of matroids $\mathcal{C}$ and a host matroid **BigM**, we define an OME as a set of matroid monomorphisms, i.e., mappings that preserve the matroid structure, from any matroid in the class $\mathcal{C}$ into **BigM** that can be constructed sequentially, only using calls to an independence oracle over the set of elements observed so far.

The use of embeddings in algorithm design is an idea that has been successfully explored in other contexts, most notably, metric embeddings both in classic settings [12, 35, 10] as well as more recently in online settings [28, 11, 38]. While our motivation and main application is the matroid secretary problem, we believe that understanding maps between matroids preserving structure is an important mathematical question in its own right that can enable other algorithmic applications beyond the matroid secretary problem.

Consequences for the MSP. We use the concept of online matroid embeddings to gain insights into the complexity of the matroid secretary problem (MSP) on *binary matroids* (see Section 2). Specifically, we relate different variants of the problem to each other and show that seemingly simpler ones are actually equivalent to harder ones (up to constant factors). See Figure 1 for an overview of the reductions that we establish in this paper.

The three variants we are interested in are: (1) the online-revealed-matroid MSP, where the matroid is a priori unknown to the online algorithm and the algorithm has access to an independence oracle on the already arrived elements, (2) the known-matroid MSP, where the structure of the matroid is known to the algorithm in advance, and (3) the prophet MSP, where the matroid structure is known in advance and additionally the weights of the elements are drawn from a known, but possibly correlated distribution.

Clearly, the online-revealed-matroid MSP is harder than the known-matroid MSP and the known-matroid MSP is harder than the prophet MSP, in the sense that if we have an α-approximation for one problem, then we also have an α-approximation for the other. Two main implications of our work are "inverses" of these statements for binary matroids, that hold up to a constant-factor loss, and apply even if we impose pairwise-independence in the prophet MSP. Such reductions between different average-case problems are notoriously difficult to achieve, as they need to ensure or maintain rather stringent assumptions on the input distribution that are essential for the target algorithm to be applicable in a meaningful way, and the required properties are easily disrupted.

Step 1: A Reduction From Online-Revealed Matroid MSP to Known-Matroid MSP. In Theorem 9 we show that the existence of an OME for a class of matroids enables a reduction from the online-revealed-matroid MSP to the known-matroid MSP. The challenge in proving this is to show that the online embedding can be used in a way that (almost) maintains uniform random arrival order. More precisely, let $\mathbf{M}$ be the unknown matroid that is revealed to the algorithm in an online fashion and let f be an OME into $\mathbf{BigM}$. Now consider the reduction: upon arrival of an element-weight pair (e, w_e) in matroid $\mathbf{M}$ at iteration t, we construct the corresponding element-weight pair $(f(e), w_e)$ as an input to MSP on $\mathbf{BigM}$ at iteration t. However, this is not a valid input to known-matroid MSP on $\mathbf{BigM}$ as it does not construct a random arrival order over the elements that are not in the image of $\mathbf{M}$.

To overcome the shortcoming of the above simple reduction, we interleave the elements in $\mathbf{BigM}$ that are in the image of $\mathbf{M}$ with the remaining elements in $\mathbf{BigM}$. In the proof of Theorem 9, our main technical argument shows that, while this interleaving does not ensure uniformity of the arrival order of the elements in $\mathbf{BigM}$, it leads to an arrival order over elements in $\mathbf{BigM}$ that is close to uniformly random arrival order in total variation distance (Section 5.1). We then complete the reduction with a coupling argument that shows that the existence of an α-competitive algorithm for the known-matroid MSP implies the existence of an $(\alpha - \epsilon)$-competitive algorithm for the online-revealed-matroid MSP.

Together with the existence of an OME that maps binary matroids into the complete binary matroid (see Section 4 and discussion below), our reduction implies that an algorithm for the MSP over binary matroids cannot meaningfully use any advance information about the matroid (Corollary 10). This is in contrast to all known $O(1)$-competitive algorithms for special cases of binary matroids [32, 16].

Step 2: A Reduction from Prophet MSP w/ Pairwise Independence to Known-Matroid MSP. In Theorem 11, we show that the existence of an OME from a class of matroids $\mathcal{C}$ to $\mathbf{BigM}$ satisfying a 2-transitivity property (that is satisfied by complete binary matroids, see definition in Section 2), allows to translate an α-competitive algorithm for prophet MSP with pairwise-independent weight distribution on $\mathbf{BigM}$ into a $C \cdot (\alpha - o(1))$-competitive algorithm for known-matroid MSP on matroid $\mathbf{M} \in \mathcal{C}$, for some constant $C > 0$.

To establish Theorem 11, we build on [18, 17] and show how to reduce prophet MSP with arbitrary correlation on $\mathbf{M} \in \mathcal{C}$ to prophet MSP with pairwise-independent weight distribution on $\mathbf{BigM}$.[1] To prove this, we first show that any uniformly random automorphism $f : \mathbf{BigM} \to \mathbf{BigM}$ satisfies the following property: for any pair of elements e, e', $\Pr[f(e) = e'] = \frac{1}{n}$ and for any two pairs of independent elements e_1, e_2 and e'_1, e'_2, $\Pr[f(e_1) = e'_1 \wedge f(e_2) = e'_2] = \frac{1}{n\cdot(n-1)}$, where $n = |\mathbf{BigM}|$ (Lemma 2).

[1] A technical detail that we are ignoring here is that our reduction is from a restricted version of the prophet MSP with arbitrary correlation, which results in an additional constant-factor loss.

This property allows us to construct an "almost pairwise independent" randomized OME $f' : \mathbf{M} \to \mathbf{BigM}$ by simply composing the given OME with a uniformly random automorphism on **BigM**. The resulting weight distribution is approximately pairwise independent in the sense that for any pair of elements $e, e' \in \mathbf{BigM}$ and weights w, w' it holds that $|\Pr[w(e) = w] \cdot \Pr[w(e') = w'] - \Pr[w(e) = w' \wedge w(e') = w']| = O\left(\frac{1}{n^2}\right)$.

To conclude the proof, we show that there exists a pairwise-independent distribution that is close to the induced weight distribution over **BigM** (Theorem 18). This is one of the most technical results of the paper (see discussion below). Combining Theorem 18 with a coupling argument similar to the one in our other reduction completes the proof.

We note that constructing an exactly k-wise independent distribution from an approximately k-wise independent distribution has been studied in previous work [4, 5, 3]. However, their techniques focus on a set of *Bernoulli* random variables with identical marginals [4, 3] or "uniformity" of the underlying random variables [5] – both conditions do not hold in our setting as the weight distribution in the prophet MSP instance can be arbitrarily correlated. In fact, in both works [4, 5], they show that if the random variables $X_1, \ldots, X_n$ satisfy "uniformity" and $|\mathbb{E}[X_i \cdot X_j] - \mathbb{E}[X_i] \cdot \mathbb{E}[X_j]| \leq \varepsilon$ then there exists pairwise independent random variables $\tilde{X}_1, \ldots, \tilde{X}_n$ within a distance of $O(n^2 \cdot \varepsilon)$ – which is not enough for our purpose as $\varepsilon = \Theta(1/n^2)$ in our case. In the following, we explain why simple approaches would tend to fail.

Technical Challenges for Step 2. One natural and simple approach could be to apply small weight perturbations to the input distribution to make $|\Pr[w(e) = w] \cdot \Pr[w(e') = w'] - \Pr[w(e) = w \wedge w(e') = w']| = o\left(\frac{1}{n^2}\right)$ and then apply results from [4, 5].

To demonstrate the technical difficulties, suppose we are given a weight distribution D over a set of elements of $\mathbf{M}$ supported over $\{w_1, \ldots, w_k\}$ such that for any $i \in [k]$, at most one element can take the weight w_i (This assumption can be made due to Lemma 16). However, we cannot ensure that each weight is definitely assigned to one of the elements of $\mathbf{M}$. Since D can be arbitrarily correlated, we can have $i, j \in [m]$, $p_i = \Pr_{w \sim D}[\exists e \in \mathbf{M} : w(e) = w_i]$ $p_{ij} = \Pr_{w \sim D}[\exists e \in \mathbf{M} : w(e) = w_i \wedge \exists e' \in \mathbf{M} : w(e') = w_j]$ with $|p_{ij} - p_i \cdot p_j| = \Omega(p_i)$. Suppose, for any $v \in \mathbf{BigM}$ and $j \in [k]$, we let random variables $X_j^v = 1$ iff $w(v) = w_j$. Now, via simple probability calculations we have

$$|\mathbb{E}[X_j^v \cdot X_j^{v'}] - \mathbb{E}[X_j^v] \cdot \mathbb{E}[X_j^{v'}]| \leq O(|p_{ij} - p_i \cdot p_j|/n^2).$$

This implies that $X_j^v : v \in \mathbf{BigM}, j \in [k]$ are almost pairwise independent. However, directly applying the results from [4, 5] only yields the existence of an exact pairwise independent distribution within distance $O(k^2 \cdot \min\{p_i, p_j\})$ from the original distribution. This is insufficient when $p_i, p_j \geq C/k^2$. To fix this, one could simply make k large by introducing new weights via small perturbations to the weight distribution. However, as such a procedure increases k, the corresponding p_i decreases, and on average, it shrinks by a factor of $1/k$. Therefore, naive approaches which apply small weight perturbations to the input distribution to increase the number of weights fail. Therefore, we have to introduce new and arguably technically involved ideas to carry out the result.

To obtain Theorem 18, we construct an explicit pairwise-independent weight distribution over **BigM** by a sequence of "small" perturbations to a naturally induced "almost" pairwise-independent distribution. At each step, we perturb the original distribution such that $\omega(1)$ many pairs of random variables end up being independent while always decreasing the pairwise correlation of the rest of the pairs. Then the main technical work is devoted to showing that the total deviation from our procedures is of the order of $\varepsilon \cdot o(n^2)$, which, combined with the fact that $\varepsilon = O(1/n^2)$ leads to the desired result. Since this construction and proof is rather technical, we defer it to the full version of the paper. We believe that our idea of sequentially constructing small perturbations would find further applications to obtain exact pairwise (or k-wise) independent distributions from their approximate counterparts in other settings.

Constructing OMEs. In Section 4 (Theorem 4 and Theorem 6), we provide a complete analysis for binary matroids. Namely, for the class of binary matroids $\mathbf{M}$ with n elements, there is an online matroid embedding into $\mathbf{BigM}$ the complete binary matroid $\mathbb{F}_2^n$. We also develop a technique for making the OME order-independent: we use properties of the automorphism group of $\mathbb{F}_2^n$ together with randomization to ensure that the images of the elements in $\mathbf{M}$ are not correlated with the arrival order. This technique is in fact more general (Theorem 8) and can be applied whenever the group of automorphisms of the host matroid is "sufficiently rich" in a sense that the theorem statement makes precise.

The key property of binary matroids that we exploit to establish these results is that in $\mathbb{F}_2^n$, there is a unique element that completes a circuit, in the sense that there cannot be two circuits of the same size that intersect in all but one element of each.

In the full version of the paper, we explore to which extent similar results can be obtained for other classes of matroids. We show that there cannot be an online matroid embedding that embeds graphic matroids into graphic matroids or, more generally, regular matroids. We also give an online matroid embedding that embeds laminar matroids into linear matroids, and show that it is impossible to embed laminar matroids into laminar matroids. Finally, we show that there is no universal host matroid that allows embedding of all matroids on n elements of a given rank. See Section 6 for an extended discussion of these results.

We believe that the lack of online matroid embeddings for graphic matroids/regular matroids that "don't leave the class" may shed light on why progress on the known-matroid MSP for graphic and regular matroids has not extended to the online-revealed version of these problems, and more generally the MSP for general binary matroids.

Approximate OMEs. In the full version of the paper, we extend the notion of an OME to allow distortion, i.e., the map approximately preserves the rank. We observe that the α-partition property in [6] and [19] can be viewed as an approximate matroid embedding into the free matroid. We thus inherit the lower bound of $\Omega(n/\log n)$ lower bound in [19] of the distortion of embedding the complete binary matroid into the free matroid.

We show tightness of this bound by constructing an online embedding which achieves a distortion of $O(n/\log n)$. We also show that there is no constant approximate online embedding of the class of graphic matroids into a free matroid when the underlying matroid is not known upfront. Therefore, any constant competitive algorithm for unknown graphic matroid secretary that relies on constructing an online embedding of the graph into a free-matroid has to exploit the random arrival order of the underlying elements or develop new techniques that do not rely on online embedding into a free matroid. A more detailed discussion appears in Section 6.

1.2 Discussion and Significance of Results

We believe that the existence or non-existence of (approximate) online matroid embeddings can shed new light on different classes of matroids and how they relate to each other. In this work, we demonstrate two implications for the matroid secretary problem.

- Our first implication (Theorem 9) offers the first formalization of the intuition that, in general, advance knowledge of the matroid structure should not help in the design of a constant-competitive algorithm for the matroid secretary problem. In light of this, it would be interesting to develop algorithms for classes of matroids for which constant-competitive algorithms exist when the algorithm has advance knowledge of the matroid structure, such as graphic matroids [32] and regular matroids [16].
- Our second implication (Theorem 11), in turn, presents a novel "line of attack" for obtaining such an algorithm for the class of binary matroids (for which no constant-competitive algorithm is known). While it was already known that it suffices to find such an algorithm for the secretary prophet version with correlated weights [18, 17], general correlated weight distributions offer little additional structure. Our result shifts the challenge away from intractable arbitrary correlations, towards the better-understood realm of pairwise independent distributions. Pairwise independent distributions admit powerful tools like concentration inequalities and have found application in areas such as hashing and constructions of pseudo-random generators (for more details, see surveys [36, 43]), and prophet inequalities [13].

1.3 Related Work

Matroid Secretary Problem. The matroid secretary problem was first studied in [9, 7, 8], who gave a $O(\log(\mathsf{rank}))$-competitive algorithm for general matroids. This bound was improved to $O(\sqrt{\log(\mathsf{rank})})$ in [14], and the state-of-the-art is a $O(\log\log(\mathsf{rank}))$-competitive algorithm [33, 22]. The algorithms of [33, 22] only uses independence oracle calls on subsets of the elements revealed so far.

For graphic matroids there is a $O(1)$-competitive algorithm, provided that the graphic matroid is known in advance [32]. The same is true for the more general class of regular matroids [16]. Laminar matroids also admit an $O(1)$-competitive algorithm [26, 29, 25]. Some evidence for the difficulty of the matroid secretary problem for general binary matroids can be found in [34] and [1], showing that binary matroids are not (b, c)-decomposable, ruling out a promising approach to obtaining an $O(1)$-competitive algorithm for this class.

Oveis Gharan and Vondrák [39] systematized the study of matroid secretary problem variants, establishing a notation for classifying problem variants according to whether the elements arrive in adversarial or random order, whether the assignment of weights to elements is adversarial or random, and whether or not the matroid structure is known in advance. In their nomenclature, the main question addressed in our work is whether the RO-AA-MK variant is equivalent to the RO-AA-MN variant for matroids in general, or for specific classes of matroids. Interestingly, for variants with adversarial arrival order but random weight assignment, [39] demonstrates a stark qualitative difference in approximability: the AO-RA-MK model (when the matroid structure is known in advance) admits a 64-competitive algorithm for all matroids, whereas the AO-RA-MN model (when the number of elements is known in advance but the matroid structure is revealed online) has no constant-competitive algorithm even for the class of *rank one matroids*!

Very recently, [41] gave a $O(1)$-competitive algorithm for the matroid secretary problem in the random assignment model when the matroid structure is not known in advance, and instead is only revealed over time. In a similar spirit, [40] presents an online contention resolution scheme for graphic matroids, that uses almost no advance information about the graph. However, they assume that the endpoints of the edges are revealed upon their arrival, which leads to an obvious OME into a graphic matroid.

Matroid Prophet Inequalities. The matroid prophet inequality problem was first studied in [24]. An asymptotically optimal $(1 - o(1))$-competitive algorithm for k-uniform matroids was given in [2]. A tight $O(1)$-competitive algorithm for the matroid prophet inequality problem was given in [31], also see [20] for the problem of maximizing submodular functions subject to matroid constraints. Constant-factor competitive algorithms can also be obtained via online contention resolution schemes (OCRS) [23]. Random-order versions of the matroid prophet inequality problem are studied in [21].

To the best of our knowledge, all these algorithms exploit that the matroid structure is known in advance. An additional difficulty for reductions of the type we present in this paper, is that typically these algorithms need to know the identity of the distribution that a certain element's weight is drawn from. For the i.i.d. case this is obviously not an obstacle, and so our reductions apply. We believe that extensions of our techniques might shed further light on the variant of the matroid prophet inequality problem, in which the matroid is revealed online.

Metric Embeddings and Distortion. An important inspiration for this work comes from the literature on metric embeddings. A classic result in this context is Bourgain's theorem [12]. The algorithmic importance of such embeddings, and Bourgain's theorem in particular, was first highlighted in the seminal papers of [35, 10].

Since then metric embeddings have found applications in a host of algorithmic problems, see, e.g., the survey of [27] and Chapter 15 of [37]. Closer to our notion of online matroid embeddings is a recent line of work on online metric embeddings [28, 11, 38] in which points of a metric space are presented one at a time to an algorithm who must then decide on a mapping to the host metric space. The main difference is that instead of preserving a matroid structure, those papers try to minimize metric distortion.

2 Matroids, Morphisms, and K-representations

Throughout the paper, we will use $[n]$ to denote the set of integers $\{1, 2, \dots, n\}$.

Matroid Definition. A matroid $\mathbf{M}$ is composed of a ground set M and a rank function $\mathsf{rank}_{\mathbf{M}} : 2^M \to \mathbb{Z}_+$ satisfying the following properties:

- $\mathsf{rank}_{\mathbf{M}}(\emptyset) = 0$;
- $\mathsf{rank}_{\mathbf{M}}(S \cup \{i\}) - \mathsf{rank}_{\mathbf{M}}(S) \in \{0, 1\}, \forall S, \{i\} \subseteq M$
- $\mathsf{rank}_{\mathbf{M}}(S \cup T) + \mathsf{rank}_{\mathbf{M}}(S \cap T) \leq \mathsf{rank}_{\mathbf{M}}(S) + \mathsf{rank}_{\mathbf{M}}(T), \forall S, T \subseteq M$ (submodularity)

It follows from the second condition that $\mathsf{rank}_{\mathbf{M}}(S) \leq |S|$. Whenever $|S| = \mathsf{rank}_{\mathbf{M}}(S)$ we say that S is an independent set of the matroid. Otherwise, we say that S is dependent. A minimal dependent set is called a *circuit*, i.e., $C \subseteq M$ is a circuit if C is dependent but every strict subset $S \subsetneq C$ is independent. We say that a matroid has rank r if $r = \max_S \mathsf{rank}_{\mathbf{M}}(S)$.

We say that an element $x \in \mathbf{M}$ is a loop if $\mathsf{rank}_{\mathbf{M}}(\{x\}) = 0$. We say that a matroid is loop-free if every set of one element is independent. Given a set $S \subseteq \mathbf{M}$ we define the span as $\mathsf{span}_{\mathbf{M}}(S) = \{x \in \mathbf{M}; \mathsf{rank}_{\mathbf{M}}(S \cup \{x\}) = \mathsf{rank}_{\mathbf{M}}(S)\}$.

Matroid Morphisms. We will use the same notation to refer to a matroid and its ground set. Given two matroids $\mathbf{M}$ and $\mathbf{N}$ we will define a morphism $f : \mathbf{M} \to \mathbf{N}$ to be a map between their ground sets that preserves rank, i.e.:

$$\mathsf{rank}_{\mathbf{N}}(f(S)) = \mathsf{rank}_{\mathbf{M}}(S), \forall S \subseteq \mathbf{M}.$$

Whenever the matroid morphism is an injective map, we will say it is a *matroid monomorphism* or a *matroid embedding*. Whenever it is bijective, we will say it is a *matroid isomorphism*. An isomorphism from a matroid to itself is called an automorphism. (Aside: this paragraph defines the category of matroids in the sense of category theory. However, we won't use any other fact from category theory other than borrowing its very convenient language.)

We refer to the set of automorphisms $\mathbf{M} \to \mathbf{M}$ as $\mathrm{Aut}(\mathbf{M})$, which forms a group under composition, i.e., given $f, g \in \mathrm{Aut}(\mathbf{M})$, then $f \circ g \in \mathrm{Aut}(\mathbf{M})$ (and $\circ$ satisfies the group axioms).

Element Copies. Given a matroid $\mathbf{M}$ and an integer k we will define the matroid $\mathbf{M}_{[k]}$ by creating k copies of each element of $\mathbf{M}$. Formally, the ground set of $\mathbf{M}_{[k]}$ is $\{(u, j); u \in \mathbf{M}, j \in [k]\}$. The rank function of $\mathbf{M}_{[k]}$ is induced by the projection $\phi : \mathbf{M}_{[k]} \to \mathbf{M}$ that maps $(u, j) \mapsto u$, i.e., $\mathsf{rank}_{\mathbf{M}_{[k]}}(S) = \mathsf{rank}_{\mathbf{M}}(\phi(S))$. By definition, the projection ϕ is a matroid morphism from $\mathbf{M}_{[k]} \to \mathbf{M}$.

If $\mathbf{N}$ is a matroid of at most n elements, every morphism $f : \mathbf{N} \to \mathbf{M}$ can be written as: $f = \phi \circ f'$ where $f' : \mathbf{N} \to \mathbf{M}_{[n]}$ is a monomorphism.

Direct Sum. Given two matroids $\mathbf{M}$ and $\mathbf{N}$, we define their direct sum $\mathbf{M} \oplus \mathbf{N}$ as the matroid whose ground set is the disjoint union of the ground sets of $\mathbf{M}$ and $\mathbf{N}$ and $\mathsf{rank}_{\mathbf{M}\oplus\mathbf{N}}(S) = \mathsf{rank}_{\mathbf{M}}(S \cap \mathbf{M}) + \mathsf{rank}_{\mathbf{N}}(S \cap \mathbf{N})$ for all S in the disjoint union of ground sets.

Graphic Matroids. We will define a few special classes of interest. We start with *graphic matroids*. Given a graph with edge set E, we can define a matroid with ground set E by defining the $\mathsf{rank}(S)$ of a subset $S \subseteq E$ as the maximum number of edges in S that don't form a cycle. We say that a matroid $\mathbf{M}$ is graphic if it is isomorphic to the matroid obtained from an undirected graph as we just described.

As an example, consider the matroid $\mathbf{M}$ with ground set $\{a, b, c\}$ and rank function such that $\mathsf{rank}(S) = |S|$. The matroid is graphic since it is isomorphic to the matroid that can be obtained from any of the graphs in Figure 2. An important thing to note, however, is that the matroid description contains no information about vertices. It only tells us which sets of edges are independent and which are not. As we can see in the figure, this is typically not enough to fully determine the graph structure.

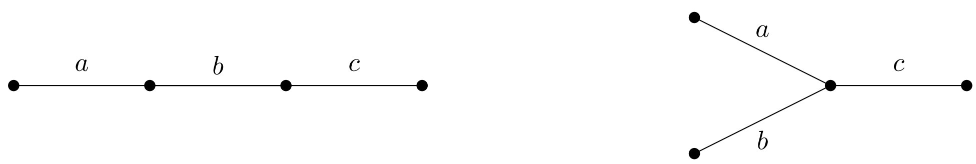

Figure 2 Two graphs that generate the same matroid on their edge set.

***K*-representable Matroids.** Let K be a field (e.g. $\mathbb{Q}, \mathbb{R}, \mathbb{F}_p$) and let K^d be the vector space formed by d-dimensional vectors with coordinates in K. We say that a subset of vectors $u_1, \ldots, u_k \in K^d$ is independent if the unique solution to $\alpha_1 u_1 + \alpha_2 u_2 + \ldots + \alpha_k u_k = 0$ for $\alpha_i \in K$ is $\alpha_1 = \alpha_2 = \ldots = \alpha_k = 0$. Any subset of K^d together with the independency relation above defines a matroid. From now on, we will use the notation K^d to represent both the vector space and the corresponding matroid. We say that a matroid $\mathbf{M}$ is K-representable if there is a matroid morphism $\mathbf{M} \to K^d$ for some integer d.

If a matroid $\mathbf{M}$ is K-representable for every field K we say that $\mathbf{M}$ is a *regular matroid*. Every graphic matroid is representable over any field by mapping an edge (u, v) to the vector $e_u - e_v$ where e_u the the u-th unit vector. (This is true even over $\mathbb{F}_2$ where $e_u - e_v = e_u + e_v$.)

For example, the matroids in Figure 2 can be represented by the vectors $(1, -1, 0, 0)$, $(0, 1, -1, 0)$, $(0, 0, 1, -1)$. As it is the case for graphic matroids, the matroid description has no information about vectors and the representation is again not unique. An equally good representation is $(1, 0, 0), (0, 1, 0), (0, 0, 1)$.

We will be specially interested in *binary matroids* which are matroids that are representable over $\mathbb{F}_2$ (the finite field of 2 elements where addition and multiplication are performed mod 2).

Laminar matroids. A family of sets, $\mathcal{A}$, is called *laminar* if it satisfies the property that for any $A, A' \in \mathcal{A}$, at least one of the sets $A \cap A'$, $A \setminus A'$, $A' \setminus A$ is empty. A *laminar matroid* $\mathbf{M}$ is one for which there exists a laminar family of sets $\mathcal{A}$ consisting of subsets of the ground set of $\mathbf{M}$ and a function $c : \mathcal{A} \to \mathbb{Z}_+$, such that the independent sets of $\mathbf{M}$ are precisely those sets $I \subseteq M$ such that $|I \cap A| \le c(A)$ for all $A \in \mathcal{A}$.

Uniform Matroid. We will denote by $\mathbf{U}_{n,r}$ the uniform matroid of n elements and rank r. This is the matroid with ground set $[n]$ and whose rank function is $\mathsf{rank}_{\mathbf{U}_{n,r}}(S) = \min(r, |S|)$. We call the $\mathbf{Fr}_n := \mathbf{U}_{n,n}$ the free matroid of rank n, i.e., a matroid of n elements in which every set is independent.

Trivial Matroid. Let $\mathbf{T}$ be the trivial matroid which has ground set $\{0\}$ and rank function $\mathsf{rank}_{\mathbf{T}}(S) = 0$ for all sets S.

2-transitive Matroid. We say that a simple matroid $\mathbf{M}$ (loop-free and no parallel elements) is 2-transitive if for any pair of independent sets of size two, $\{e_1, e'_1\}, \{e_2, e'_2\} \in \mathbf{M}$ there exists an automorphism $f \in \mathrm{Aut}(\mathbf{M})$ satisfying $f(e_1) = e'_1$ and $f(e_2) = e'_2$. There are several matroids that satisfies the 2-transitive property including complete affine matroids, complete projective matroids, free matroids, and their truncations [30].

3 Online Matroid Embeddings

We are interested in studying matroids whose structure is revealed to an algorithm in an online fashion. For that, it will be useful to take into account the order in which elements are processed, which we will represent by an indexing of the ground set: $\pi : [n] \to \mathbf{M}$.

A matroid $\mathbf{N}$ is a restriction of matroid $\mathbf{M}$ if the ground set of $\mathbf{N}$ is a subset of the ground set of $\mathbf{M}$ and $\mathsf{rank}_{\mathbf{M}}$ coincides with $\mathsf{rank}_{\mathbf{N}}$ on the ground set of $\mathbf{N}$.

Given a matroid with ordered ground set specified by a pair $(\mathbf{M}, \pi)$, we say that $(\mathbf{M}', \pi')$ is a prefix-restriction of $(\mathbf{M}, \pi)$ if $n' = |\mathbf{M}'| < |\mathbf{M}|$, $\mathbf{M}'$ is the restriction of $\mathbf{M}$ to $\pi([n'])$ and π' is the restriction of π to $[n']$.

Let $\mathcal{C}$ be a class of matroids that is closed under restriction (e.g., the class of all matroids, graphic matroids, binary matroids, K-representable matroids, matroids of rank at most r). An *online matroid morphism* (OMM) for class $\mathcal{C}$ consists of a host matroid $\mathbf{BigM}$, together with matroid morphisms

$$f_{\mathbf{M},\pi} : \mathbf{M} \to \mathbf{BigM}$$

for every $\mathbf{M} \in \mathcal{C}$ and every indexing $\pi : [n] \to \mathbf{M}$ of the ground set of $\mathbf{M}$, such that for every prefix-restriction $(\mathbf{M}', \pi')$ of $(\mathbf{M}, \pi)$, the map $f_{\mathbf{M}',\pi'}$ is the restriction of $f_{\mathbf{M},\pi}$ to the ground set of $\mathbf{M}'$.

If all morphisms $f_{\mathbf{M},\pi}$ are monomorphisms, we say that they form an *online matroid embedding* (OME). Given an online matroid morphism it is easy to construct an online matroid embedding by copying the elements of $\mathbf{BigM}$.

▶ **Lemma 1.** *Let $\mathcal{C}$ be a class of matroids, where each matroid $\boldsymbol{M} \in \mathcal{C}$ has at most n elements and $f_{\boldsymbol{M},\pi} : \boldsymbol{M} \to \boldsymbol{BigM}$ form an online matroid morphism. Then there is an online matroid embedding $f'_{\boldsymbol{M},\pi} : \boldsymbol{M} \to \boldsymbol{BigM}_{[n]}$.*

Proof. We define $f'_{\mathbf{M},\pi}$ as follows: for each $u \in \mathbf{M}$ if $u = \pi(k)$ let $f'_{\mathbf{M},\pi}(u) = (f_{\mathbf{M},\pi}(u), j)$ where $j = |i \in [k]; f_{\mathbf{M},\pi}(\pi(i)) = u|$. The functions $f'_{\mathbf{M},\pi}$ are injective by construction and they are matroid morphisms by the definition of $\mathbf{BigM}_{[n]}$. In fact: $f_{\mathbf{M},\pi} = \phi \circ f'_{\mathbf{M},\pi}$ where ϕ is the natural projection $\mathbf{BigM}_{[n]} \to \mathbf{BigM}$. Finally note that they can be constructed online since the identity of the copy used is only a function of the set of elements that arrived up to this point. ◀

With this definition we can ensure that an online algorithm is able to construct a monomorphism from an unknown matroid $\mathbf{M}$ to $\mathbf{BigM}$ in an online fashion. Consider a matroid $\mathbf{M}$ for which the elements arrive according to π. At each time t, we can observe the structure of the matroid $\mathbf{M}_t$ which is the restriction of $\mathbf{M}$ to $\pi([t])$. Let π_t be the restriction of π to $[t]$. If we have an online matroid embedding, we can first construct $f_{\mathbf{M}_1,\pi_1}$, then extend to $f_{\mathbf{M}_2,\pi_2}$ and so forth.

It will also be convenient to define a randomized online matroid morphism (embedding) which for every matroid $\mathbf{M} \in \mathcal{C}$ and ordering π specifies a distribution over (mono)morphisms $f_{\mathbf{M},\pi} : \mathbf{M} \to \mathbf{BigM}$ such that for every prefix-restriction $(\mathbf{M}', \pi')$ the distribution of the restriction of $f_{\mathbf{M},\pi}$ to the ground set of $\mathbf{M}'$ coincides with the distribution of $f_{\mathbf{M}',\pi'}$.

Finally, we say that a randomized online matroid embedding is order-independent if the distribution of $f_{\mathbf{M},\pi}$ doesn't depend on π. In other words, for any two orderings π and π', the morphisms $f_{\mathbf{M},\pi}$ and $f_{\mathbf{M},\pi'}$ are equally distributed.

Uniform Order-Independent Embedding. Let f be an order-independent online embedding from $\mathbf{M} \to \mathbf{BigM}$. We consider an order independent-randomized embedding $g : \mathbf{M} \to \mathbf{BigM}$ by composing f with uniformly random automorphism $f' \in \mathrm{Aut}(\mathbf{BigM})$, i.e. $g = f' \circ f$.

Interestingly, whenever $\mathbf{BigM}$ satisfies 2-transitive property then the embedding g maps each element of $e \in \mathbf{M}$ uniformly at random over the matroid $\mathbf{BigM}$. In addition, for any independent set of pair of elements $\{e, e'\} \subseteq \mathbf{M}$ and pair of elements $\{\tilde{e}, \tilde{e}'\} \subseteq \mathbf{BigM}$, the events $\{g(e) = \tilde{e}\}$ and $\{g(e') = \tilde{e}'\}$ are "almost" independent. More formally,

▶ **Lemma 2.** *Given a simple host matroid (loop-free and no parallel elements) $\boldsymbol{BigM}$ such that for any two pairs of distinct elements $\{e_1, e_2\}$ and $\{e'_1, e'_2\}$ there exists an automorphism $f' \in \mathrm{Aut}(\boldsymbol{BigM})$ satisfying $f'(e_1) = e'_1$ and $f'(e_2) = e'_2$, then a uniformly random automorphism f sampled from $\mathrm{Aut}(\boldsymbol{BigM})$ satisfies:*

1. *For any $e, e' \in \boldsymbol{BigM}$, $\Pr[f(e) = e'] = \frac{1}{n}$.*
2. *For any two pairs of elements e_1, e_2 and e'_1, e'_2 (s.t. $e_1 \neq e_2$ and $e'_1 \neq e'_2$), we have*

$$\Pr[f(e_1) = e'_1 \wedge f(e_2) = e'_2] = \frac{1}{n \cdot (n-1)}$$

Proof. Consider the action of the group of automorphisms $\mathrm{Aut}(\mathbf{BigM})$ on the set of pairs of distinct elements $P = \{(e_1, e_2) : e_1, e_2 \in \mathbf{BigM}, e_1 \neq e_2\}$. For two pairs $(e_1, e_2), (e'_1, e'_2) \in P$, the set of automorphisms $f \in \mathrm{Aut}(\mathbf{BigM})$ that satisfy $f(e_1, e_2) = (e'_1, e'_2)$ is nonempty, and therefore, it is a coset of the subgroup of stabilizers of (e_1, e_2) (i.e., the set of automorphisms that satisfy $f(e_1, e_2) = (e_1, e_2)$). Since all cosets of a subgroup must have the same size, and because all pairs $(e'_1, e'_2) \in P$ define a different coset, for a uniformly drawn automorphism f,

$$\Pr[f(e_1, e_2) = (e'_1, e'_2)] = \frac{1}{|P|} = \frac{1}{n \cdot (n-1)}.$$

An analogous argument gives that $\Pr[f(e) = e'] = 1/n$. ◀

Online vs. Offline Embeddings. The difficulty of constructing an online matroid embedding is that the elements of **BigM** corresponding to certain elements of **M** must be chosen before the full matroid structure of **M** is known. If we merely wanted to construct a matroid **BigM** that contains an isomorphic copy of every matroid in $\mathcal{C}$, that would be very easy: **BigM** could be taken to be the direct sum of all the matroids in $\mathcal{C}$.

4 OMEs for Binary Matroids

Before we discuss how to use online matroid embeddings in online algorithms, it is important to show first that they exist in non-trivial cases. For that, we will provide a complete analysis for binary matroids. Recall that $\mathbb{F}_2^n$ is the complete binary matroid of rank n and that graphic matroids and regular matroids are special cases of binary matroids. Our first result is the existence of an OME for this class. This will be done by showing the existence of an OMM and using Lemma 1 to convert an OMM to an OME.

Our first step is to show a lemma that the matroid $\mathbb{F}_2^n$ is special in the sense that its group of matroid automorphisms coincides with its group of vector space automorphisms:

▶ **Lemma 3.** *A mapping $A : \mathbb{F}_2^n \to \mathbb{F}_2^n$ is a matroid automorphism iff it is an automorphism of vector spaces.*

Proof. An automorphism of vector spaces $A : \mathbb{F}_2^n \to \mathbb{F}_2^n$ is a bijection such that for any vectors $v_1, \ldots, v_k \in \mathbb{F}_2^n$ it holds that $A(\sum_{i=1}^k v_i) = \sum_{i=1}^k Av_i$. This in particular implies that a set of vectors $v_1, \ldots, v_k$ is independent iff the vectors $Av_1, \ldots, Av_k$ are independent. This is because a coefficient vector $(\alpha_1, \ldots, \alpha_k)$ satisfies the equation $\sum_{i=1}^k \alpha_i v_i = 0$ if and only if it satisfies $\sum_{i=1}^k \alpha_i A v_i = A(\sum_{i=1}^k \alpha_i v_i) = 0$, so the first equation has only trivial solutions if and only if the second equation has only trivial solutions.

For the opposite direction, if A is a matroid automorphism and $e_1, \ldots, e_n$ is the standard basis of $\mathbb{F}_2^n$ then $Ae_1, \ldots, Ae_n$ must be linearly independent elements of $\mathbb{F}_2^n$. Now, take any vector $v = \sum_{i \in S} e_i$. Since $\{v\} \cup \{e_i; i \in S\}$ forms a circuit, then $\{Av\} \cup \{Ae_i; i \in S\}$ must form a circuit. Since the only non-zero constant in $\mathbb{F}_2$ is 1, it must hold that: $Av + \sum_{i \in S} Ae_i = 0$ and hence $Av = \sum_{i \in S} Ae_i$. Hence A is also an automorphism of vector spaces. ◀

▶ **Theorem 4.** *Let $\mathcal{C}$ be the class of binary matroids of at most n elements and let **BigM** be the complete binary matroid $\mathbb{F}_2^n$. Then there exists an OMM for $\mathcal{C}$ into **BigM**.*

Proof. Given a binary matroid **M** we construct a mapping $f : \mathbf{M} \to \mathbb{F}_2^n$ as follows. We keep a counter k initially set to 1. For each element a we process, if it is independent of the previously arrived elements (i.e. there are no circuits containing a and the elements seen so far), we set $f(a) = e_k$ and increment k. Otherwise, a forms a circuit with a set of previously arrived elements $u_1, \ldots, u_m$ for some integer $m \geq 0$. This means that their image $f(a), f(u_1), \ldots, f(u_m)$ must be a minimal $\mathbb{F}_2$-linearly dependent set. Since the only non-zero constant in $\mathbb{F}_2^n$ is 1, then it must hold that:

$$f(a) + f(u_1) + \ldots + f(u_m) = 0$$

and hence we can map: $f(a)$ to $f(u_1) + \ldots + f(u_m)$ (recall that $1 = -1$ in $\mathbb{F}_2$).

Finally, we need to argue that f is a matroid morphism. Observe that if $\mathbf{M}$ is a binary matroid, then there exists a morphism $g : \mathbf{M} \to \mathbb{F}_2^n$. Let $\{b_1, \ldots, b_r\}$ be the elements of $\mathbf{M}$ such that $f(b_i) = e_i$. By the fact that g is matroid morphism, $g(b_1), \ldots, g(b_r)$ are linearly independent elements in $\mathbb{F}_2^n$. By Lemma 3 there is an automorphism $A \in \mathrm{Aut}(\mathbb{F}_2^n)$ that takes $g(b_i)$ to e_i. Since matroid morphisms compose, $Ag : \mathbf{M} \to \mathbb{F}_2^n$ is matroid morphism.

$$\begin{tikzcd} \mathbf{M} \arrow[r, "f"] \arrow[dr, "g"'] & \mathbb{F}_2^n \\ & \mathbb{F}_2^n \arrow[u, "A"'] \end{tikzcd}$$

Finally, we argue that $f(a) = Ag(a)$ for all a in $\mathbf{M}$. We show this by induction. For each element processed by the algorithm, if it is independent from previously arrived elements, then $f(a) = Ag(a)$ by construction. Otherwise, there are previously arrived elements such that $a, u_1, \ldots, u_m$ form a circuit. Hence their image under Ag must be a linearly dependent set of $\mathbb{F}_2$-vectors, which means that:

$$Ag(a) = Ag(u_1) + \ldots + Ag(u_m) = f(u_1) + \ldots + f(u_m) = f(a)$$

where the second equality holds by induction. Since f coincides with Ag, f is a matroid morphism. ◀

Furthermore, there is a randomized OME that is order-independent. We will show it as a consequence of the following lemmas:

▶ **Lemma 5.** *Given a binary matroid $\boldsymbol{M}$ and the complete binary matroid $\mathbb{F}_2^n$, if there are two matroid morphisms $f, g : \boldsymbol{M} \to \mathbb{F}_2^n$, then there exist an automorphism $A \in Aut(\mathbb{F}_2^n)$ such that $f = A \circ g$.*

Proof. Let r be the rank of $\mathbf{M}$ and let $\{b_1, \ldots, b_r\}$ be a basis of $\mathbf{M}$. Then $\{f(b_1), f(b_2), \ldots, f(b_r)\}$ and $\{g(b_1), g(b_2), \ldots, g(b_r)\}$ are both sets of independent vectors in $\mathbb{F}_2^n$. Then there exists an automorphism A of vector spaces (and hence a matroid automorphism) that sends $g(b_i)$ to $f(b_i)$. For any other element in $v \in \mathbf{M}$ consider any circuit formed with a subset of the basis. If $\{v\} \cup \{b_i; i \in S\}$ is a circuit then it must be the case that: $f(v) = \sum_{i \in S} f(b_i)$ and $g(v) = \sum_{i \in S} g(b_i)$. Given that A is a automorphism of vector spaces, then: $Ag(v) = \sum_{i \in S} Ag(b_i) = \sum_{i \in S} f(b_i) = f(v)$. ◀

▶ **Theorem 6.** *There is a order-independent randomized OMM from the class of binary matroids $\mathcal{C}$ into the complete binary matroid.*

Proof. Let $f_{\mathbf{M},\pi}$ be the online matroid morphism constructed in Theorem 4 and consider $A \circ f_{\mathbf{M},\pi}$ when A is drawn uniformly at random from $\mathrm{Aut}(\mathbb{F}_2^n)$. It is clear that for every fixed A the morphisms $A \circ f_{\mathbf{M},\pi}$ still form an OMM. We only need to check that they are order-independent. To see that, observe that if π and π' are two different orderings of the ground set of $\mathbf{M}$ then $f_{\mathbf{M},\pi}$ and $f_{\mathbf{M},\pi'}$ are two morphisms $\mathbf{M} \to \mathbb{F}_2^n$. By the previous lemma, there is $A_0 \in \mathrm{Aut}(\mathbb{F}_2^n)$ such that $f_{\mathbf{M},\pi} = A_0 \circ f_{\mathbf{M},\pi'}$. Now, the distribution of $A \circ f_{\mathbf{M},\pi}$ for a random $A \sim \mathrm{Aut}(\mathbb{F}_2^n)$ is the same distribution as $A \circ A_0 \circ f_{\mathbf{M},\pi'}$ which is the same distribution of $A \circ f_{\mathbf{M},\pi'}$, since $A_0 \circ A$ is also uniformly distributed over $\mathrm{Aut}(\mathbb{F}_2^n)$. ◀

Extending to Copies. In the following section we will be needing an online matroid embedding. For that reason, we need to extend the last two theorems to deal with copies. The extension is rather simple: we only need to observe that a matroid autormorphism of the commplete binary matroid with n copies of each element $(\mathbb{F}_2^n)_{[n]}$ can be decomposed into an automorphism $A \in \mathrm{Aut}(\mathbb{F}_2^n)$ and indexings of the identities of the copies.

▶ **Lemma 7.** *If $f \in Aut((\mathbb{F}_2^n)_{[n]})$ then there exists $A \in Aut(\mathbb{F}_2^n)$ and indexings $\sigma_u : [n] \to [n]$ for each $u \in \mathbb{F}_2^n$ such that $f((u, j)) = (Au, \sigma_u(j))$*

Proof. Let id be the identity map and $\phi : (\mathbb{F}_2^n)_{[n]} \to \mathbb{F}_2^n$ the natural projection. Now, $\phi \circ f$ and ϕ are two matroid morphisms from $(\mathbb{F}_2^n)_{[n]} \to \mathbb{F}_2^n$ so by Theorem 5 there is $A \in \mathrm{Aut}(\mathbb{F}_2^n)$ such that $\phi \circ f = A \circ \phi$ (see the commutative diagram below). This means in particular that $f((u, j)) = (Au, \sigma_u(j))$ for some indexings σ_u.

$$\begin{tikzcd}
& (\mathbb{F}_2^n)_{[n]} \arrow[r, "\phi"] & \mathbb{F}_2^n \\
(\mathbb{F}_2^n)_{[n]} \arrow[ur, "f"] \arrow[dr, "id"'] & & \\
& (\mathbb{F}_2^n)_{[n]} \arrow[r, "\phi"] & \mathbb{F}_2^n \arrow[uu, "A"']
\end{tikzcd}$$

◀

With that, Theorem 6 automatically extends to the matroid with copies $(\mathbb{F}_2^n)_{[n]}$ by taking a random automorphism from group $\mathrm{Aut}((\mathbb{F}_2^n)_{[n]})$.

Remark A (Single Orbit Morphisms). Theorem 6 is not particular to binary matroids. The only fact it uses is that all the morphisms of the OMM belong to the same orbit under the action of the automorphism group $\mathrm{Aut}(\mathbf{BigM})$. We can state it more generally as follows. The proof is identical to Theorem 6, so we omit it here.

▶ **Theorem 8** (Generalization of Theorem 6)**.** *Let $f_{\mathbf{M},\pi}$ be an OMM of class $\mathcal{C}$ into $\boldsymbol{BigM}$ such that given two orderings π and π' of the ground set of $\boldsymbol{M}$, there is an automorphism $A \in Aut(\boldsymbol{BigM})$ such that $f_{\mathbf{M},\pi'} = A \circ f_{\mathbf{M},\pi}$. Then there is an order independent randomized OMM of class $\mathcal{C}$ into $\boldsymbol{BigM}$.*

Remark B (Other Fields). In this section we repeatedly use the fact that $\mathbb{F}_2$ has only one non-zero constant, so whenever we identify a circuit in the matroid, we know exactly what is the linear dependency between the elements in the corresponding vector field. This is no longer true even in slightly larger fields like $\mathbb{F}_3$. If vectors $u, v, w \in \mathbb{F}_3^n$ form a circuit, it could be that: $w = \pm u \pm v$ in the representation. As a consequence, given two matroid morphisms $f, g : \mathbf{M} \to \mathbb{F}_3^n$ there may not exist a vector-space automorphism A of $\mathbb{F}_3^n$ such that $f = Ag$. In the previous example, if $f(u) = g(u)$, $f(v) = g(v)$ but $f(w) = f(u) + f(v)$ but $g(w) = g(u) - g(v)$ no such automorphism can exist.

5 OMEs and the Matroid Secretary Problem

In this section, we use OMEs to explore the complexity of the matroid secretary problem (MSP) on binary matroids. We consider three versions of the problem, each making different assumptions on the data generation process and what's known to the algorithm. We will use OMEs to establish equivalences between these problems, showing that seemingly simpler problems are actually equivalent to harder ones (up to constants).

Three Versions of the MSP. We consider the following three versions of the MSP, and aim to establish the relations in Figure 1. In all three variants, the goal is an algorithm for selecting elements that form an independent set, and whose combined weight is in expectation an α-approximation to the weight of the optimal basis.

- **Online-revealed-matroid MSP:** In this version of the problem, there is an underlying matroid $\mathbf{M}$, which is a priori unknown to the algorithm. The algorithm has only access to the number of elements $n = |\mathbf{M}|$ and to a promise that $\mathbf{M} \in \mathcal{C}$ for a class of matroids $\mathcal{C}$. For each element $u \in \mathbf{M}$, an adversary determines a weight $w_u \in \mathbb{R}_+$. The algorithm processes pairs (u, w_u) in random order at each time, but it only knows the rank function restricted to the subset of elements that have already arrived. Upon seeing the element, the algorithm must irrevocably decide whether to accept that element or not, subject to the constraint that the set of accepted elements must be an independent set of $\mathbf{M}$.
- **Known-matroid MSP:** In this version, the matroid $\mathbf{M}$ is known to the algorithm ahead of time. The only information missing is the weight of each element, which is again chosen adversarially. As before elements arrive in random order, and the algorithm must make immediate accept/reject decisions, with the restriction that the chosen set of elements must be an independent set of $\mathbf{M}$.
- **Prophet MSP:** In this version, the matroid $\mathbf{M}$ is again known to the algorithm ahead of time. However, this time the weight of each element is drawn from a known distribution $\mathcal{D}$, which can potentially sample weights in a correlated manner. As in the other versions elements are then presented to the algorithm in random order, and the algorithm aims to select an independent set of high weight in an online manner.

Note that the first version is clearly harder than the second and the second version is clearly harder than the third, in the sense that an α-approximation to the harder problem immediately implies an α-approximation to the simpler one. We derive approximate "inverses" of these comparisons from the existence of OMEs, even if we restrict prophet MSP to pairwise-independent distributions.

Our Reductions. We first use OMEs to show an (essentially exact) "inverse" of the comparison between the online-revealed-matroid MSP and the known-matroid MSP, implying that for binary matroids the latter is as hard as the former.

▶ **Theorem 9.** *If a class $\mathcal{C}$ of matroids admits a randomized order-independent online matroid embedding into matroid $\boldsymbol{BigM}$, then an α-approximation to the known-matroid MSP for $\boldsymbol{BigM}_{[n]}$ implies that for every $\epsilon > 0$ there is a $(\alpha - \epsilon)$-approximation to the online-revealed-matroid MSP for $\mathcal{C}$.*

For the case of binary matroids that we previously discussed, the matroids $\mathbf{BigM}$ and $\mathbf{BigM}_{[n]}$ themselves are binary in which case we can obtain the following corollary:

▶ **Corollary 10.** *For binary matroids, there is no gap in approximability between the known-matroid MSP and the online-revealed-matroid MSP.*

As our second result, we use the existence of OMEs into a 2-transitive host matroid as a tool to establish the approximate equivalence of known-matroid MSP and prophet MSP with pairwise-independent distributions.

▶ **Theorem 11.** *Suppose a class $\mathcal{C}$ of matroids admits a randomized order-independent online matroid embedding into matroid **BigM** which is* 2*-transitive. Then an* α*-approximation to the prophet MSP with pairwise-independent weight distributions for* $\mathbf{BigM}_{[k]}$*, implies that for some constant* $C > 0$ *there is a* $C \cdot (\alpha - o(1))$*-approximation to the known-matroid MSP for* $\mathcal{C}$*.*

Noting that the full binary matroid is 2-transitive, and there exists an order-independent OME from the class of binary matroids into a full binary matroid, we obtain the following corollary.

▶ **Corollary 12.** *For binary matroids, there is a constant-factor gap in approximability between the known-matroid MSP and the prophet MSP with pairwise-independent weight distributions.*

We note that we can also chain the two reductions, and this way relate the online-revealed-matroid MSP to the Prophet MSP with pairwise-independent distributions. The rest of this section is devoted to the proofs of the reductions.

5.1 Proof of Theorem 9

Let $\mathbf{M}$ be the unknown matroid that is revealed to the algorithm in an online fashion and let $f_{\mathbf{M},\pi}$ be an order-independent randomized OME into $\mathbf{BigM}$. Let $n = |\mathbf{M}|$, $N = |\mathbf{BigM}|$ and d, k be two larger integers (to be specified later) where k is a multiple of d.

An instance of the MSP consists of a sequence of weighted elements from $\mathbf{M}$ that are presented to the algorithm in random order. Our goal is to map it on the fly to a random instance of the MSP on $\mathbf{BigM}_{[k]}$. The main difficulty is, as usual, doing it online and preserving the random order. Our strategy will be to first provide an *offline reduction* which preserves random order and obtains the desired approximation, but can't be implemented online. After that we will provide a *mostly-online implementation* of this reduction, i.e. a procedure that samples from the same distribution generated by the offline reduction and that with $1 - \epsilon$ probability can be implemented online. With the remaining ϵ probability, the process *raises a flag*. Raising a flag will indicate that from that point on, the reduction can no longer be implemented online. Algorithmically, we will stop the algorithm whenever we raise a flag and obtain zero reward. Finally, we will show that the probability of raising a flag is very small for large values of k and d.

Offline Reduction. We will view a weighted element of $\mathbf{M}$ as a pair (u, w_u) with $u \in \mathbf{M}$ and $w_u \in \mathbb{R}_+$. In the offline reduction, we assume we have access to the entire matroid $\mathbf{M}$ and the entire sequence of weights. Now, we will produce a distribution of instances of $\mathbf{BigM}_{[k]}$ as follows.

For each element v in $\mathbf{BigM}$ sample k different i.i.d. timestamps t_{vj} for $j \in [k]$ from the Uniform$([0, 1])$ distribution. Those timestamps specify the arrival time of each of the k copies of the elements in $\mathbf{BigM}$ and induce a random ordering over the ground set of $\mathbf{BigM}_{[k]}$. For the matroid $\mathbf{M}$, sample a random embedding $f : \mathbf{M} \to \mathbf{BigM}$ from the OME. (Since the embedding is order independent, we don't need to know the arrival order of elements in $\mathbf{M}$ to sample such embedding). For each $u \in \mathbf{M}$, pick a random copy of $f(u)$ and set its weight to w_u. For the remaining elements, set the weight equal to zero.

This random input is clearly in random order as it is equivalent to starting with k copies of the elements of $\mathbf{BigM}$ where all but one copy has weight zero if that corresponds to an element of $\mathbf{M}$ and randomly permuting those elements. Now, feed this instance to the

α-competitive algorithm for the MSP on $\mathbf{BigM}_{[k]}$. From the set selected by the algorithm, discard any element with zero weight chosen by the algorithm. The elements with non-zero weight chosen in $\mathbf{BigM}$ correspond to an independent set of $\mathbf{M}$ with the same weight. Hence, in expectation, we select an α-approximation to the optimal basis of $\mathbf{M}$.

Mostly-online Implementation. The drawback of the previous reduction is that it can't be implemented online as we are assuming we know everything in advance. We will describe the same sampling procedure in a way that with high probability we can generate the instance as we go. In the sampling procedure, we will also define an event *raise a flag* which will mean that we can't generate that instance online as we learn the structure of the matroid $\mathbf{M}$.

The process will again start by sampling i.i.d. timestamps t_{vj} for $v \in \mathbf{BigM}$ and $j \in [k]$ from Uniform$([0, 1])$. In addition, we will also sample n additional timestamps from Uniform$([0, 1])$ sort them in increasing order and denote the sorted list by $T_1, \ldots, T_n$.

We will now divide the interval $[0, 1]$ into intervals $I_i = [\frac{i-1}{d}, \frac{i}{d})$ for $i \in [d]$. If more than one timestamp T_s falls in the same interval I_i, we will raise a flag. We will count how many of the timestamps $t_{v,j}$ fall in each interval:

$$X_{vi} = |\{j \in [k]; t_{vj} \in I_i\}|$$

With that, also define:

$$A_{vi} = \min\left(\frac{k}{d}, X_{vi}\right) \qquad B_{vi} = \max\left(0, X_{vi} - \frac{k}{d}\right)$$

Now, process the elements of the matroid $\mathbf{M}$ according to order π (which will be sampled at random). As we process the s-th element $u = \pi(s) \in \mathbf{M}$, we will map it to an element in $v = f_{\mathbf{M},\pi}(u) \in \mathbf{BigM}$ using the randomized OME. Now, we will apply the following procedure to choose a copy of v in $\mathbf{BigM}_{[k]}$ to assign weight w_u:

- find the interval I_i containing T_s.
- with probability $A_{vi}d/k$, choose one of the X_{vi} timestamps t_{vj} in interval I_i
- with remaining probability (if any), raise a flag and choose a different interval $I_{i'}$ with probability proportional to $B_{vi'}$ and choose a timestamp t_{vj} in that interval.

We assign weight w_u to the element with the chosen timestamp and zero weight to others. Now, we will show the following facts.

▶ **Lemma 13.** *The* mostly online implementation *samples sequences with the same probability as the* offline reduction*.*

Proof. Observe that if we ignore the weights, the order of the elements of $\mathbf{BigM}$ is the same in both processes since they are determined by the timestamps t_{vj}. What we are left to argue is that for each v we select uniformly random timestamp t_{vj} to assign the non-negative weight. For that, observe that since $f_{\mathbf{M},\pi}$ is order independent, it has the same distribution as if we first sampled a monomorphism $f : \mathbf{M} \to \mathbf{BigM}$, and then we sampled an independent uniform indexing π for the arrival order of the elements in $f(\mathbf{M})$. This implies that when we assign the timestamps $T_1, \ldots, T_n$ according to π, the resulting distribution is the same as if we assigned i.i.d. Uniform$[0, 1]$ timestamps T_v to each element $v \in f(\mathbf{M})$, and therefore, the interval T_v lands in is uniformly chosen and independent across elements $v \in f(\mathbf{M})$. Now fix a certain timestamp t_{vj} and let I_i be the interval containing it. We will show that the probability that this timestamp is selected is exactly $1/k$.

Consider two cases: either $X_{vi} \le k/d$ in which case the probability of sampling t_{vj} is the probability that the timestamp T_v is in I_i (which is $1/d$), times the probability we decide to sample a timestamp inside I_i (which is $X_{vi}d/k$), times the probability that out of those, we choose t_{vj} (which is $1/X_{vi}$). The total probability is:

$$\frac{1}{d} \cdot \frac{X_{vi}d}{k} \cdot \frac{1}{X_{vi}} = \frac{1}{k}$$

In the case where $X_{vi} > k/d$, then it is possible that we sample t_{vj} also when T_v is outside I_i. The probability that we sample t_{vj} and T_v is in I_i is:

$$\frac{1}{d} \cdot 1 \cdot \frac{1}{X_{vi}} = \frac{1}{dX_{iv}}$$

The probability that we sample when it is outside is the probability that we choose a different interval $I_{i'}$, raise a flag and then move to interval I_i, which is:

$$\sum_{i'} \frac{1}{d}\left(1 - \frac{A_{vi'}d}{k}\right) \cdot \frac{B_{vi}}{\sum_{i'} B_{vi'}} \cdot \frac{1}{X_{vi}} = \frac{1}{k}\frac{(k - \sum_{i'} A_{vi'})}{\sum_{i'} B_{vi'}} \cdot \frac{B_{vi}}{X_{vi}} = \frac{1}{k}\frac{B_{vi}}{X_{vi}}$$

because $\sum_{i'} A_{vi'} + \sum_{i'} B_{vi'} = \sum_{i'} X_{vi'} = k$. Taking those two probabilities together, we have:

$$\frac{1}{dX_{iv}} + \frac{1}{k}\frac{B_{vi}}{X_{vi}} = \frac{1}{k}.$$ ◀

▶ **Lemma 14.** *If no flags were raised, we can produce the instance on the fly as we process $\boldsymbol{M}$.*

Proof. Let $i_1 < \ldots < i_n$ be the indices of the intervals such that $T_s \in I_{i_s}$. Since no flag was raised, then each T_s landed in a different interval and the s-th element that arrives from matroid $\mathbf{M}$ is mapped to a copy inside I_{i_s}. This enables the following online reduction: once the s-th element arrives we can decide the weights of all the elements in intervals $I_{i_{s-1}+1}$ to I_{i_s} and feed to the MSP algorithm for $\mathbf{BigM}_{[k]}$. In this sub-sequence there will be at most one element on non-zero weight which corresponds to the arriving element of $\mathbf{M}$. We can observe if that element was selected in $\mathbf{BigM}_{[k]}$ and if so, we can select it in $\mathbf{M}$. ◀

▶ **Lemma 15.** *For any n and ϵ, there are large enough k and d, such that the probability that we raise a flag is at most ϵ.*

Proof. The first event in which we raise a flag is when two timestamps T_s land in the same interval. The probability that this happens is at most n^2/d. Now, note that for each interval i and each of the n elements v in $\mathbf{BigM}$ that have non-zero weights, we have by the Chernoff bound that:

$$\mathbb{P}\left(X_{vi} \le (1-\delta)\frac{k}{d}\right) \le \exp\left(-\frac{\delta^2 k}{2d}\right)$$

Hence with probability at most $nd \exp\left(-\frac{\delta^2 k}{2d}\right)$, the timestamps t_{vj} are such that the probability we raise a flag when we try to choose a timestamp in the same interval as T_s is more than δn. Taking the union bound of those events, we get:

$$\frac{n^2}{d} + nd\exp\left(-\frac{\delta^2 k}{2d}\right) + \delta n$$

Taking $\delta = \epsilon/(3n)$, $d = 3n^2/\epsilon$ and k large enough, we get that the total probability of raising a flag is at most ϵ. ◀

Taking those lemmas together, we can conclude the proof of Theorem 9. For that, let Alg be an α-competitive algorithm for $\mathbf{BigM}_{[k]}$ and let Y represent the sequence of the MSP sampled by the offline reduction. Let's represent by $\mathsf{Alg}(Y)$ the weight of the elements selected by Alg and Opt the weight of the optimal basis. By the fact that the offline reduction produces an instance in random order, we know that $\mathbb{E}[\mathsf{Alg}(?'Y)] \geq \alpha\mathsf{Opt}$.

Our online reduction, will attempt to construct Y on the fly. If we raise the flag, we will stop the algorithm and pretend we had zero reward. If not, we will continue the reduction and collect $\mathsf{Alg}(Y)$ reward. We will denote by Flag the event that the flag was raised and by $\overline{\mathsf{Flag}}$ its complement. Our total reward will be:

$$\mathbb{E}[\mathsf{Alg}(Y)\cdot\mathbf{1}\{\overline{\mathsf{Flag}}\}] = \mathbb{E}[\mathsf{Alg}(Y)]-\mathbb{E}[\mathsf{Alg}(Y)\cdot\mathbf{1}\{\mathsf{Flag}\}] \geq \mathbb{E}[\mathsf{Alg}(Y)]-\mathsf{Opt}\cdot\mathbb{P}[\mathsf{Flag}] \geq (\alpha-\epsilon)\mathsf{Opt}.$$

5.2 Proof of Theorem 11

Let $\mathbf{M}$ be the unknown matroid with $|\mathbf{M}| = n$ that is revealed to the algorithm in an online fashion, which admits an order-independent randomized OME into $\mathbf{BigM}$ with $|\mathbf{BigM}| = M$. First, we obtain the following simple reduction that allows us to focus on a special class of the prophet MSP in which each element takes a weight from the set of weights W with $|W| = m = O(n^2)$. In addition, we can restrict the weight distribution such that each element has a distinct weight from the weight class.

▶ **Lemma 16.** *If there exists an α-approximation to the prophet MSP on $\boldsymbol{M}$ with weight distribution $\mathcal{D}$ supported over the set of weights W with $\mathsf{rank}(\boldsymbol{M}) = d$, $|\boldsymbol{M}| = n$, $|W| = O(n^2)$ with $\max_{w\in W} w \leq 1$ and $\mathbb{E}_{\mathcal{D}}[\boldsymbol{OPT}(\boldsymbol{M})] \in \left[\frac{1}{16}, 1\right]$ such that for all $w \in W$, there exists at most one element assigned weight of w with probability one, then there exists an $\left(\frac{\alpha}{256} - \frac{1}{32d}\right)$-approximation to prophet MSP on matroid $\boldsymbol{M}$ with any arbitrary weight distribution.*

The proof of the above lemma simply follows from Sublemma-4.2 from [18] that reduces any arbitrary prophet MSP with $O(\log(|\mathsf{rank}(\mathbf{M})|))$ many weights and $\mathbb{E}_{\mathcal{D}}[\mathbf{OPT}(\mathbf{M})] \in \left[\frac{1}{16}, 1\right]$. We then add distinct noise of the order of $O\left(\frac{1}{n^2}\right)$ to ensure that the weight of each element is distinct. We defer the details of this reduction to the full version.

For simplicity, we let $W = \{w_1, \ldots, w_m\}$ and consider the prophet MSP on matroid $\mathbf{M}$ and weight distribution $\mathcal{D}$ supported over the set of weights W satisfying the conditions from Lemma 16.

Extending BigM with Copies. We let $\mathbf{BigM}_{[m\cdot N]}$ be a matroid with $m \cdot N$ parallel copies of each element of $\mathbf{BigM}$ with $|\mathbf{BigM}| = M$ and integer $N = \Omega\left(2^{M^2}\right)$. We divide the set of $N \cdot m$ copies into m sets of size N, each part corresponding to weight class w_i. We use N_i to denote the set of labels corresponding to weight w_i for all $i \in [m]$ with $|N_i| = N$. We sometimes denote N_i by $[N] = \{1, 2, \ldots, N\}$ whenever it is clear from the context which w_i we are referring to. In addition, the ℓ-th copy of the weight class corresponding to weight w_i of element $\mathbf{v} \in \mathbf{BigM}$ is denoted as $\mathbf{v}^{i,\ell}$.

Reduction to "Almost" Pairwise Independent Prophet MSP. We first define the weight distribution $\mathcal{D}^*$ over $\mathbf{BigM}_{[m\cdot N]}$ in Definition 17, which is "almost" pairwise independent. Then in Theorem 18, we show an existence of exact pairwise independent weight distribution $\tilde{\mathcal{D}}$ over $\mathbf{BigM}_{[m\cdot N]}$ which is "close" to the distribution defined in Definition 17 in total-variation distance. This allows us to utilize the fact that any algorithm $\mathcal{A}$ can not distinguish between the almost pairwise independent weight distribution $\mathcal{D}^*$ and $\tilde{\mathcal{D}}$ with high probability. Finally, we complete the proof of Theorem 11.

We proceed with the formal proof. We begin by defining the almost pairwise independent weight distribution over $\mathbf{BigM}_{[m\cdot N]}$.

▶ **Definition 17** (Almost P.W. Independent Distribution)**.** *Consider the weight distribution $\mathcal{D}^*$ over the elements of $\boldsymbol{BigM}_{[m\cdot N]}$ defined as follows:*

1. *Given an order independent OMM $f' : \boldsymbol{M} \to \boldsymbol{BigM}$, we sample a random automorphism $f'' \in \mathrm{Aut}(\boldsymbol{BigM})$ and obtain an order independent matroid morphism $f = f'' \circ f'$.*
2. *For any $\mathbf{v} \in \boldsymbol{M}$ with $w(\mathbf{v}) = w_i$, let $\mathbf{u} = f(\mathbf{v})$. We sample $\ell \sim \mathrm{Unif}(N_i)$ and assign the weight of $w(\mathbf{u}^{i,\ell}) = w_i$.*
3. *We assign the weight of the rest of the elements of $\boldsymbol{BigM}_{[m\cdot N]}$ to be zero.*

We first observe that when N is much larger than M, the distribution in the above definition is almost pairwise independent. To see this, we first observe that for any $i \in [m], \ell \in N_i$ and $\mathbf{u} \in \mathbf{BigM}$, $\mathbf{u}^{i,\ell} \in \mathbf{BigM}_{[m\cdot N]}$ can potentially either take a weight of w_i or zero. For simplicity, now consider any two distinct elements $\mathbf{u}_{i,\ell}, \mathbf{u}'_{j,\ell'} \in \mathbf{BigM}_{[m\cdot N]}$ and weight distribution $\mathcal{D}$ such that there always exists a pair of elements $\mathbf{v}, \mathbf{v}'' \in \mathbf{M}$ that are assigned weights of w_i, w_j, respectively. Since f' is a random automorphism $\mathrm{Aut}(\mathbf{BigM})$, we have $\Pr[w(\mathbf{u}_{i,\ell}) = w_i] = \Pr[f(\mathbf{v}) = \mathbf{u}] \cdot \frac{1}{N} = \frac{1}{M\cdot N}$. On the other hand, we have

$$\Pr[w(\mathbf{u}_{i,\ell}) = w_i \wedge w(\mathbf{u}_{j,\ell}) = w_j] = \Pr[f'(\mathbf{v}) = \mathbf{v} \wedge f'(\mathbf{v}') = \mathbf{u}'] \cdot \frac{1}{N^2} = \frac{1}{M(M-1)} \cdot \frac{1}{N^2},$$

which is close to the product $\Pr[w(\mathbf{u}_{i,\ell}) = w_i] \cdot \Pr[w(\mathbf{u}_{j,\ell}) = w_j] = \frac{1}{M^2\cdot N^2}$. However, in general, we can not guarantee that $\mathcal{D}$ will always assign weights w_i, w_j to some pair of elements of $\mathbf{v}, \mathbf{v}'$ of $\mathbf{M}$. In addition, the above argument also fails if we have $i = j$ or $\mathbf{u} = \mathbf{u}'$ as in both of these cases, $\Pr[w(\mathbf{u}_{i,\ell}) = w_i \wedge w(\mathbf{u}_{j,\ell}) = w_j] = 0$. Intuitively, we can circumvent these pairwise correlation issues by taking N large enough as it makes pairwise correlations small enough. More precisely,

$$|\Pr[w(\mathbf{u}_{i,\ell}) = w_i \wedge w(\mathbf{u}_{j,\ell}) = w_j] - \Pr[w(\mathbf{u}_{i,\ell}) = w_i] \cdot \Pr[w(\mathbf{u}_{j,\ell}) = w_j]| = O\left(\frac{1}{M^2 \cdot N^2}\right).$$

Using this observation, we prove the following technical theorem.

▶ **Theorem 18.** *For $N = \Omega\left(2^{M^2}\right)$, let weight distribution $\mathcal{D}^*$ over $\boldsymbol{BigM}_{[m\cdot N]}$ be defined as in Definition 17, then there exist a pairwise-independent weight distribution $\tilde{\mathcal{D}}$ over $\boldsymbol{BigM}_{[m\cdot N]}$ such that $\mathrm{TV}_{\mathcal{D}^*,\tilde{\mathcal{D}}} \leq O\left(\frac{m^3}{M}\right)$.*

The proof of the above theorem is highly technical and deferred to the full version. The idea is to construct an explicit $\tilde{\mathcal{D}}$ by a sequence of small perturbations to $\mathcal{D}^*$. We emphasize that the choice of $N = \Omega(2^{M^2})$ is required due to the limitations of the techniques developed to prove Theorem 18. We conjecture that one can prove the similar theorem for $N = \Omega(\mathrm{Poly}(M))$, which we leave as an intriguing technical open problem.

Proof of Theorem 11. To prove the main theorem, we first prove the following: if there exists an α-approximate algorithm to the prophet MSP instance with a pairwise independent weight distribution for the matroid $\mathbf{BigM}_{[m\cdot N]}$, then there exists an $(\alpha - o(1))$-approximate algorithm for the prophet MSP instance for the matroid $\mathbf{M}$ with weight distribution $\mathcal{D}$ supported over the set of weights W with $\mathsf{rank}(\mathbf{M}) = d$, $|\mathbf{M}| = n$ and $|W| = O(n^2)$ with $\max_{w\in W} w \leq 1$ and $\mathbb{E}_{\mathcal{D}}[\mathbf{OPT}(\mathbf{M})] \in \left[\frac{1}{16}, 1\right]$ such that for all $w \in W$, there exists at most one element assigned weight of w with probability one. Combining this with Lemma 16, we will conclude the proof of the Theorem 11.

Given an order-independent matroid morphism $f' : \mathbf{M} \to \mathbf{BigM}$, we sample a random automorphism $f'' \in \mathrm{Aut}(\mathbf{BigM})$ and obtain an order-independent matroid morphism $\tilde{f} = f'' \circ f'$. Given $\tilde{f} : \mathbf{M} \to \mathbf{BigM}$, we obtain $f : \mathbf{M} \to \mathbf{BigM}_{[m]}$ that maps each $\mathbf{v} \in \mathbf{M}$ to $\mathbf{u}^i \in \mathbf{BigM}_{[m]}$ iff $\tilde{f}(\mathbf{v}) = \mathbf{u}$ and $w(\mathbf{v}) = w_i$. We then consider $\mathbf{BigM}_{[m\cdot N]}$, i.e. matroid $\mathbf{BigM}_{[m]}$ with N many copies of each element.

We first consider an offline reduction as follows: for all $\mathbf{v}^i \in \mathbf{BigM}_{[m]}$, sample N many independent arrival times from $\mathrm{Unif}[0,1]$ denoting the uniformly random arrival times of elements of $\mathbf{BigM}_{[m\cdot N]}$. Given any $\mathbf{v} \in \mathbf{M}$ and pair $(\mathbf{v}, w(\mathbf{v}) = w_i)$, with $\mathbf{u}^i = f(\mathbf{v})$, let $w(\mathbf{u}^{i,\ell}) = w_i$ uniformly random from $\ell \in [N]$.

Since f does not require any information about the arrival order, the above-described offline reduction is a valid instance of matroid prophet secretary over $\mathbf{BigM}_{[m\cdot N]}$. In addition, the induced weight assignment over $\mathbf{BigM}$ due to offline reduction is identical to the distribution $\mathcal{D}^*$ defined in Definition 17.

Given the uniformly random arrival of elements of $\mathbf{M}$, we construct an "almost online implementation" of the above offline reduction similar to the proof of Theorem 9. We let N be large enough $(\Omega(n^2/\varepsilon))$ such that the probability of "almost online implementation" raising a flag is at most ε. In fact, for the proof of Theorem 18, we let $N = \Omega(2^{M^2})$ which satisfies the required condition.

We let $\mathcal{A}$ be an α-approximate algorithm for the pairwise-independent prophet MSP on matroid $\mathbf{BigM}_{[m\cdot N]}$. Since there exists a pairwise independent weight distribution $\tilde{\mathcal{D}}$ over $\mathbf{BigM}$ within the total variation distance of $O\left(\frac{m^3}{M}\right)$, the algorithm $\mathcal{A}$ can not distinguish the weight distribution $\mathcal{D}^*$ from $\tilde{\mathcal{D}}$ with probability at least $1 - O\left(\frac{m^2}{M}\right)$.

Let $\mathcal{E}$ be the event when the algorithm $\mathcal{A}$ can not distinguish between $\mathcal{D}^*$ and $\tilde{D}$. We note that when event $\mathcal{E}$ does not hold, the offline optimal can be bounded by

$$\mathbb{E}[\mathbf{OPT}(\mathbf{BigM}_{m\cdot N}) \mid \mathcal{E}^c] \le \mathsf{rank}(\mathbf{BigM}_{[m\cdot N]}) \cdot \max_{w_i \in W} w_i \le n \cdot 1 = n,$$

where the second inequality follows because $\mathsf{rank}(\mathbf{BigM}_{[m\cdot N]}) = n$ and $w_i \le 1$ for all $i \in [m \cdot N]$.

Now, let S be the selected set of elements of $\mathbf{BigM}$ by $\mathcal{A}$ w.r.t. weight distribution $\mathcal{D}^*$. We can bound,

$$\begin{aligned}\mathbb{E}[w(S)] &= \mathbb{E}[w(S) \mid \mathcal{E}] \cdot \Pr[\mathcal{E}] + \mathbb{E}[w(S) \mid \mathcal{E}^c] \cdot \Pr[\mathcal{E}^c] \\ &\le \mathbb{E}[w(S) \mid \mathcal{E}] + \mathbb{E}[\mathbf{OPT}(\mathbf{BigM}_{[M\cdot m]}) \mid \mathcal{E}^c] \cdot \frac{m^3}{M} \\ &\le \mathbb{E}[w(S) \mid \mathcal{E}] + \frac{m^3 \cdot n}{M}.\end{aligned}$$

Above, the first inequality holds because $\Pr[\mathcal{E}^c] \le \frac{m^2}{M}$ and the second inequality holds because $\mathbb{E}[\mathbf{OPT}(\mathbf{BigM}_{[m\cdot N]}) \mid \mathcal{E}^c] \le n$. Due to our reduction, the performance of the algorithm on the original Prophet MSP instance $\mathbb{F}'$ is lower bounded by $\mathbb{E}[w(S) \mid \mathcal{E}]$, next we lower bound the expectation $\mathbb{E}[w(S) \mid \mathcal{E}]$,

$$\begin{aligned}\mathbb{E}[w(S) \mid \mathcal{E}] &\ge \mathbb{E}[w(S)] - \frac{m^3 \cdot n}{M} \\ &\ge \alpha \cdot \mathbb{E}_{\tilde{\mathcal{D}}}[\mathbf{OPT}(\mathbf{BigM}_{[m\cdot N]})] - o(1) \cdot \mathbf{OPT}(\mathbf{M}) \\ &\ge \mathbf{OPT}(\mathbf{M}) \cdot \left(\alpha - \frac{\alpha \cdot m^3}{M}\right) - o(1) \cdot \mathbf{OPT}(\mathbf{M}) \ge (\alpha - o(1)) \cdot \mathbf{OPT}(\mathbf{M}).\end{aligned}$$

Above, the second inequality holds because $\mathcal{A}$ is an α approximate algorithm for Prophet MSP with pairwise independent prior and $\frac{m^3 \cdot n}{M} = O(m^4/2^m) = o(1)$. The third inequality holds because $\mathrm{TV}_{\mathcal{D}^*,\tilde{\mathcal{D}}} \leq O(m^3/M)$. Finally, since the performance of the reduction on the original matroid secretary is

$$\mathbb{E}[w(S) \mid \mathcal{E}] \cdot \Pr[\text{Reduction does not Flag}] \geq (1-\varepsilon) \cdot (\alpha - o(1)) \cdot \mathbf{OPT}(\mathbf{M}).$$

This concludes the proof. ◀

6 Additional Results

OMEs Beyond Binary Matroids. Our main result from Section 4 is an online matroid embedding that embeds binary matroids into binary matroids. We refer to online matroid embeddings where both $\mathbf{M}$ and $\mathbf{BigM}$ are of the same class as "within-class" OMEs. In the full version of the paper, we explore whether such within-class OMEs can exist for graphic matroids. We show that such embeddings cannot exist, and in fact, we show that graphic matroids cannot be embedded in an online-fashion to regular matroids. To rule out the existence of such an online embedding, we show that if it would exist, then $\mathbf{BigM}$ must contain an isomorphic copy of $\mathbb{F}_2^n$. However, $\mathbb{F}_2^n$ contains an isomorphic copy of the Fano plane which is not representable over $\mathbb{F}_3$ [42]. Hence $\mathbf{BigM}$ can't be regular.

We then we give another example of an OME, namely for laminar matroids. We show how to embed the class of laminar matroids $\mathbf{M}$ with at most n elements into $\mathbf{BigM}$ which is a complete linear matroid of rank n over any field with sufficiently many elements. On the flip side, we show that there is no OME if we require $\mathbf{BigM}$ to also be laminar.

Finally, we show an impossibility result of constructing an OME for the class of all matroids. This is shown by studying finite projective planes and showing that for those matroids, elements that haven't arrived yet impose non-trivial constraints on the already arrived elements. As a corollary we obtain an impossibility of constructing an OME for the class of all matroids representable over fields of characteristic at least 7.

Approximate Embeddings. Motivated by these negative results, in the full version we also define and study matroid embeddings that allow some degree of *distortion*. We define a β-approximate embedding $f : \mathbf{M} \to \mathbf{N}$ as a map between ground sets that approximately preserves rank, i.e.,

$$\frac{1}{\beta} \cdot \mathsf{rank}_{\mathbf{M}}(S) \leq \mathsf{rank}_{\mathbf{N}}(f(S)) \leq \mathsf{rank}_{\mathbf{M}}(S), \forall S \subseteq \mathbf{M},$$

together with corresponding randomized and online versions of this notion.

By connecting the α-partition property of [1, 19] to our notion of β-approximate embedding, we obtain a lower bound of $\Omega(n/\log n)$ on the distortion β that any randomized offline embedding from a complete binary matroid into a free matroid must incur. We complement this lower bound with a matching upper bound of $O(n/\log n)$, and show that this randomized embedding can be computed in an online manner.

Finally, we explore approximate embeddings for graphic matroids. We translate a result of [32] to our formalism, obtaining a 2-approximate offline embedding from graphic matroids into a free matroid. This embedding requires knowledge of the precise graphical structure, in particular, whether two edges are adjacent to the same node. We show that there is no constant-approximate online embedding when the elements are revealed online, and the algorithm only has access to an independence oracle.

References

1 Dorna Abdolazimi, Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. Matroid partition property and the secretary problem. In *ITCS 2023*, pages 2:1–2:9, 2023. doi:10.4230/LIPIcs.ITCS.2023.2.

2 Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM J. Comput.*, 43(2):930–972, 2014. doi:10.1137/120878422.

3 Noga Alon, Alexandr Andoni, Tali Kaufman, Kevin Matulef, Ronitt Rubinfeld, and Ning Xie. Testing k-wise and almost k-wise independence. In *STOC 2007*, pages 496–505, 2007. doi:10.1145/1250790.1250863.

4 Noga Alon, Oded Goldreich, and Yishay Mansour. Almost k-wise independence versus k-wise independence. *Inf. Process. Lett.*, 88(3):107–110, 2003. doi:10.1016/S0020-0190(03)00359-4.

5 Noga Alon and Shachar Lovett. Almost k-wise vs. k-wise independent permutations, and uniformity for general group actions. In *APPROX'12*, pages 350–361, 2012. doi:10.1007/978-3-642-32512-0_30.

6 Moshe Babaioff, Michael Dinitz, Anupam Gupta, Nicole Immorlica, and Kunal Talwar. Secretary problems: weights and discounts. In *Proceedings of the twentieth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1245–1254. SIAM, 2009. doi:10.1137/1.9781611973068.135.

7 Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. A knapsack secretary problem with applications. In *APPROX-RANDOM 2007*, pages 16–28, 2007. doi:10.1007/978-3-540-74208-1_2.

8 Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Matroid secretary problems. *J. ACM*, 65(6):35:1–35:26, 2018. doi:10.1145/3212512.

9 Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA 2007*, pages 434–443, 2007. URL: http://dl.acm.org/citation.cfm?id=1283383.1283429.

10 Yair Bartal. On approximating arbitrary metrices by tree metrics. In *STOC 1998*, pages 161–168, 1998. doi:10.1145/276698.276725.

11 Yair Bartal, Nova Fandina, and Seeun William Umboh. Online probabilistic metric embedding: A general framework for bypassing inherent bounds. In *SODA 2020*, pages 1538–1557, 2020. doi:10.1137/1.9781611975994.95.

12 Jean Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Isr. J. Math.*, 52:46–52, 1985.

13 Ioannis Caragiannis, Nick Gravin, Pinyan Lu, and Zihe Wang. Relaxing the independence assumption in sequential posted pricing, prophet inequality, and random bipartite matching. In *WINE 2022*, pages 131–148, 2022.

14 S. Chakraborty and O. Lachish. Improved competitive ratio for the matroid secretary problem. In *SODA 2012*, pages 1702–1712, 2012.

15 Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In Leonard J. Schulman, editor, *STOC 2010*, pages 311–320, 2010. doi:10.1145/1806689.1806733.

16 Michael Dinitz and Guy Kortsarz. Matroid secretary for regular and decomposable matroids. *SIAM J. Comput.*, 43(5):1807–1830, 2014. doi:10.1137/13094030X.

17 Shaddin Dughmi. The outer limits of contention resolution on matroids and connections to the secretary problem. In *ICALP 2020*, volume 168 of *LIPIcs*, pages 42:1–42:18, 2020. doi:10.4230/LIPIcs.ICALP.2020.42.

18 Shaddin Dughmi. Matroid secretary is equivalent to contention resolution. In *ITCS 2022*, volume 215 of *LIPIcs*, pages 58:1–58:23, 2021. doi:10.4230/LIPIcs.ITCS.2022.58.

19 Shaddin Dughmi, Yusuf Hakan Kalayci, and Neel Patel. Limitations of stochastic selection problems with pairwise independent priors. In *STOC 2024*, pages 479–490, 2024. doi:10.1145/3618260.3649718.

20 Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM J. Comput.*, 49(3):540–582, 2020. doi:10.1137/20M1323850.

21 Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. In *SODA 2018*, pages 700–714, 2018. doi:10.1137/1.9781611975031.46.

22 Moran Feldman, Ola Svensson, and Rico Zenklusen. A simple $O(\log\log(\text{rank}))$-competitive algorithm for the matroid secretary problem. *Math. Oper. Res.*, 43(2):638–650, 2018. doi:10.1287/MOOR.2017.0876.

23 Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes with applications to bayesian selection problems. *SIAM J. Comput.*, 50(2):255–300, 2021. doi:10.1137/18M1226130.

24 Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI 2007*, pages 58–65, 2007. URL: http://www.aaai.org/Library/AAAI/2007/aaai07-009.php.

25 Zhiyi Huang, Zahra Parsaeian, and Zixuan Zhu. Laminar matroid secretary: Greedy strikes back. *arXiv preprint arXiv:2308.09880*, 2023. doi:10.48550/arXiv.2308.09880.

26 Sungjin Im and Yajun Wang. Secretary problems: Laminar matroid and interval scheduling. In *SODA 2011*, pages 1265–1274, 2011. doi:10.1137/1.9781611973082.96.

27 Piotr Indyk. Algorithmic applications of low-distortion geometric embeddings. In *FOCS 2001*, pages 10–33, 2001. doi:10.1109/SFCS.2001.959878.

28 Piotr Indyk, Avner Magen, Anastasios Sidiropoulos, and Anastasios Zouzias. Online embeddings. In *APPROX-RANDOM 2010*, pages 246–259, 2010. doi:10.1007/978-3-642-15369-3_19.

29 Patrick Jaillet, José A Soto, and Rico Zenklusen. Advances on matroid secretary problems: Free order model and laminar case. In *IPCO 2013*, pages 254–265, 2013. doi:10.1007/978-3-642-36694-9_22.

30 William M Kantor. Homogeneous designs and geometric lattices. *J. Comb. Theory Ser. A*, 38(1):66–74, 1985. doi:10.1016/0097-3165(85)90022-6.

31 Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In *STOC 2012*, pages 123–136, 2012. doi:10.1145/2213977.2213991.

32 Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs. In *ICALP 2009*, pages 508–520, 2009. doi:10.1007/978-3-642-02930-1_42.

33 Oded Lachish. O(log log rank) competitive ratio for the matroid secretary problem. In *FOCS 2014*, pages 326–335, 2014. doi:10.1109/FOCS.2014.42.

34 Marilena Leichter, Benjamin Moseley, and Kirk Pruhs. On the impossibility of decomposing binary matroids. *Oper. Res. Lett.*, 50(5):623–625, 2022. doi:10.1016/J.ORL.2022.09.003.

35 Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15:215–245, 1995. doi:10.1007/BF01200757.

36 Michael Luby, Avi Wigderson, et al. Pairwise independence and derandomization. *Foundations and Trends® in Theoretical Computer Science*, 1(4):237–301, 2006.

37 Jiri Matousek. *Lectures on Discrete Geometry*. Springer, New York, NY, 2002.

38 Ilan Newman and Yuri Rabinovich. Online embedding of metrics. *Isr. J. Math.*, 2025. Forthcoming.

39 Shayan Oveis Gharan and Jan Vondrák. On variants of the matroid secretary problem. *Algorithmica*, 67:472–497, 2013. doi:10.1007/S00453-013-9795-Y.

40 Richard Santiago, Ivan Sergeev, and Rico Zenklusen. Simple random order contention resolution for graphic matroids with almost no prior information. In *SOSA 2023*, pages 84–95, 2023. doi:10.1137/1.9781611977585.CH8.

41 Richard Santiago, Ivan Sergeev, and Rico Zenklusen. Constant-competitiveness for random assignment matroid secretary without knowing the matroid. *Math. Program.*, 210(1):815–846, 2025. doi:10.1007/S10107-024-02177-X.

42 William Thomas Tutte. A homotopy theorem for matroids. I, II. *Trans. Am. Math. Soc.*, 88(1):144–174, 1958.

43 Salil P. Vadhan. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012. doi:10.1561/0400000010.