A Match Made by Entropy

Mohammad Reza Aminian, Rad Niazadeh, Pranav Nuti

Abstract

Online contention resolution schemes (OCRSs) are a basic tool for converting ex-ante feasible fractional solutions to relaxations of problems into online policies that remain feasible in every realized sample path. They play a central role in prophet inequalities, Bayesian online allocation, and online mechanism design. In this note, we provide a particularly simple algorithm and analysis for the bipartite matching OCRS problem. This algorithm is conjectured to be optimal.

1 Introduction

A recurring strategy in Bayesian online decision making problems is to first compute an *ex-ante feasible* plan that satisfies constraints only in expectation, and then to convert it into an online policy that preserves *ex-post* feasibility in every realization. OCRSs are a technical tool that provide a principled way to perform this conversion. The general OCRS framework for downward-closed feasibility systems was developed in [6], and has since become a standard black-box rounding primitive for Bayesian selection problems.

In the bipartite matching online contention resolution scheme problem, we have a bipartite graph $G = (U \cup V, E)$, and each edge is independently active with known probability x_e, where $x = (x_e)_{e \in E}$ is a fractional matching. The edges arrive online in an arbitrary order (chosen potentially by an adversary with no knowledge of the activations and any random bits the algorithm uses), and an algorithm must (irrevocably) select a matching.

An algorithm is said to be α-selectable if

$$\mathbb{P}[e \text{ is selected} \mid e \text{ is active}] \geq \alpha \quad \text{for every edge } e.$$

The goal is to make α as large as possible. An algorithm with a large selectability immediately provides several applications; for instance, an α-selectable OCRS yields an α competitive ratio for the bipartite matching prophet inequality (see [10] for a discussion of the relationship between OCRSs and prophet inequalities).

For bipartite matching prophet inequalities, the simple nonadaptive pricing policy of Gravin and Wang [7] gives a $1/3$ guarantee. Later work improved this to 0.337 [5] and then to 0.349 [13]. These improvements are technically involved.

Ma, MacRury, and Nuti [11] identified $(3-\sqrt{5})/2 \approx 0.382$ as the natural "independence benchmark" for matching OCRSs and left open whether it could be achieved. They also showed that this factor could not be beaten by any OCRS satisfying a certain natural concentration property, and that no OCRS could obtain a selectability larger than 0.390.

The constant $(3 - \sqrt{5})/2$ is a root of the equation $\alpha = (1 - \alpha)^2$, which explains why it is the independence benchmark: if an edge with a tiny probability of activation arrives last, and its two endpoints are available roughly independently (with probability $1 - \alpha$ each since those endpoints were already selected with probability α), then the edge can be selected, conditional on its activation, with probability at most $(1 - \alpha)^2$, so $\alpha \leq (1 - \alpha)^2$.

Our main result in this note, an exposition of our result in [2], gives exactly this selectability for bipartite graphs.

Theorem 1. *For every bipartite graph $G = (U \cup V, E)$, and every activation vector $x \in [0,1]^E$ that is a fractional matching, there is an OCRS that selects a matching and is $(3-\sqrt{5})/2$-selectable.*

In a departure from previous work, our algorithm uses sampling from *Gibbs distributions* on matchings. For most of the analysis, we will not actually use that the graph is bipartite. The only fact we will need about bipartite graphs (not true for general graphs) is that when we sample a matching from a Gibbs distribution, the events that two given vertices on different sides of the partition are unmatched are positively correlated, as we show in Section B. Exploiting this correlation significantly simplifies our analysis in comparison to [5, 11, 13].

2 The algorithm and analysis

Let $G = (U \cup V, E)$ be a bipartite graph and let $\mathcal{M}$ be the family of matchings of G. A fractional matching is a vector $x \in [0,1]^E$ satisfying

$$\sum_{e \in \delta(z)} x_e \le 1 \qquad \forall z \in U \cup V,$$

where $\delta(z)$ denotes the set of edges incident to a vertex z. We assume for convenience that $x_e > 0$ for every edge; edges with $x_e = 0$ can simply be ignored. Set

$$\alpha := \frac{3-\sqrt{5}}{2}, \qquad p_e := \alpha x_e.$$

Let μ be the maximum-entropy distribution over matchings with marginals p, i.e.,

$$\mathbb{P}_{M\sim\mu}[e \in M] = p_e \qquad \forall e \in E.$$

As recalled in Section A, this distribution exists and has a Gibbs form: there are weights $w_e > 0$ such that

$$\mu(M) = \frac{1}{Z(w)} \prod_{e \in M} w_e, \qquad Z(w) = \sum_{M \in \mathcal{M}} \prod_{e \in M} w_e.$$

Furthermore, these weights can be computed approximately in polynomial time [14, 15]. For each edge, define

$$\rho_e := \frac{w_e}{1+w_e}.$$

To start the algorithm, sample a random matching $R \sim \mu$. (This can be done approximately in polynomial time; see Section A.) The set R should be interpreted as "real selected edges from the past plus predictions for the future." When an edge e arrives, we first delete its old prediction from R. If the edge is active and can be added to the remaining matching ($R \setminus \{e\}$), we accept it and add it into R with probability ρ_e/x_e. For now, let us assume that the maximum-entropy distribution satisfies $\rho_e \le x_e$, meaning ρ_e/x_e is an actual probability.

The reason this acceptance probability rule is natural is that it ensures that the distribution of R remains unchanged as the algorithm runs, as the following proposition demonstrates:

Proposition 1. *Suppose that $\rho_e \le x_e$ $\forall e \in E$. Then Algorithm 1 is well-defined and, after every arrival, R continues to have distribution μ. Consequently, the final output has distribution μ.*

Algorithm 1 Simulate-then-replace for bipartite matching

Require: fractional matching x, Gibbs weights w with marginals $p = \alpha x$

Sample $R \sim \mu$, the maximum-entropy distribution over matchings with marginals p
$\triangleright$ R is the real selected matching plus future predictions
for each arriving edge $e = (u, v)$ **do**
    $R \leftarrow R \setminus \{e\}$ $\triangleright$ forget the old prediction for e
    if e is active and $R \cup \{e\} \in \mathcal{M}$ **then**
        With probability ρ_e/x_e, accept e and set $R \leftarrow R \cup \{e\}$
        Otherwise reject e
    else
        Reject e
return R $\triangleright$ at the end, all predictions have been replaced by online decisions

Proof. Consider the initially sampled set R. Fix an edge e and a matching $T \subseteq E \setminus \{e\}$. Conditioned on $R \setminus \{e\} = T$, the only possible matchings are T and, if feasible, $T \cup \{e\}$. Their Gibbs weights differ by the multiplicative factor w_e. Hence

$$\mathbb{P}_{R \sim \mu}[e \in R \mid R \setminus \{e\} = T] = \begin{cases} \rho_e, & T \cup \{e\} \in \mathcal{M}, \\ 0, & T \cup \{e\} \notin \mathcal{M}. \end{cases}$$

We now prove that the distribution of R always remains μ by induction over the arrival sequence. Initially, $R \sim \mu$ by construction. Suppose that immediately before edge e arrives, we have $R \sim \mu$, and let $T = R \setminus \{e\}$. The activation of e is independent of T and occurs with probability x_e. If $T \cup \{e\}$ is feasible, Algorithm 1 inserts e with probability

$$x_e \cdot \frac{\rho_e}{x_e} = \rho_e.$$

If $T \cup \{e\}$ is infeasible, it never inserts e. Therefore, conditional on $R \setminus \{e\} = T$, the post-update value of the e-th coordinate is distributed exactly as it is under μ conditional on $R \setminus \{e\} = T$. All other edges remain equal to T. Thus the post-update matching again has law μ.

By induction, $R \sim \mu$ after every arrival. At the end of the process, every initial prediction has been deleted (when its edge arrived) and has been replaced by the corresponding online decision. Hence the returned set is the set of accepted active edges and has distribution μ. $\square$

The proposition implies that the final output has distribution μ. Since μ has marginals $p_e = \alpha x_e$, and since the algorithm only selects active edges, we get

$$\mathbb{P}[e \text{ is selected} \mid e \text{ is active}] = \frac{\mathbb{P}[e \text{ is selected}]}{x_e} = \frac{\alpha x_e}{x_e} = \alpha.$$

It remains to prove that the inequalities $\rho_e \le x_e$ do indeed hold.

Proposition 2. *For the maximum-entropy distribution μ defined above with marginals $p_e = \alpha x_e$, we have*

$$\rho_e \le x_e \qquad \forall e \in E.$$

Proof. Fix an edge $e = (u, v)$ with $u \in U$ and $v \in V$. Sample M from μ, and define the addability event

$$\mathsf{Add}(e) := \{(M \setminus \{e\}) \cup \{e\} \in \mathcal{M}\}.$$

By the Gibbs calculation used in the proof of Proposition 1,

$$p_e = \mathbb{P}[e \in M] = \sum_T \mathbb{P}_{M\sim\mu}[e \in M \mid M \setminus \{e\} = T]\mathbb{P}_{M\sim\mu}[M \setminus \{e\} = T] = \rho_e \, \mathbb{P}[\mathsf{Add}(e)].$$

Therefore, since $p_e = \alpha x_e$, it is enough to prove

$$\mathbb{P}[\mathsf{Add}(e)] \geq \alpha.$$

For a vertex z, let $\mathsf{Matched}(z)$ be the event that z is matched by M. Since M is always a matching and the edge marginals of μ are $p_f = \alpha x_f$,

$$\mathbb{P}[\mathsf{Matched}(z)] = \sum_{f \in \delta(z)} \mathbb{P}[f \in M] = \alpha \sum_{f \in \delta(z)} x_f \leq \alpha.$$

Therefore, the probability z is unmatched is:

$$\mathbb{P}[\overline{\mathsf{Matched}(z)}] \geq 1 - \alpha.$$

If both u and v are unmatched, then e is addable. Now for every $u \in U$ and $v \in V$, the events that u is unmatched and that v is unmatched are positively correlated. This positive-correlation statement is well-known for matchings sampled from Gibbs distributions [8], but for the sake of completeness, we provide a proof in Section B. Thus,

$$\begin{aligned}
\mathbb{P}[\mathsf{Add}(e)] &\geq \mathbb{P}[\overline{\mathsf{Matched}(u)} \cap \overline{\mathsf{Matched}(v)}] \\
&\geq \mathbb{P}[\overline{\mathsf{Matched}(u)}]\mathbb{P}[\overline{\mathsf{Matched}(v)}] \\
&\geq (1-\alpha)^2 = \alpha,
\end{aligned}$$

where the last equality follows from $\alpha = (3-\sqrt{5})/2$. This is the desired bound for every edge e. □

Combining Proposition 1 and Proposition 2 completes the proof of the theorem stated in the introduction.

3 Further discussion

The algorithm in this note is one instance of a more general design strategy from [2]. In the bipartite matching setting, we choose a random matching from a maximum-entropy Gibbs distribution, and then maintain this distribution throughout the online process. When an edge arrives, the algorithm deletes its prediction for that edge and replaces it with a real online decision using the Gibbs weights. The analysis reduces to a question about when $\rho_e \leq x_e$, or in other words, for a random matching drawn from a Gibbs distribution, how likely is a given edge to be addable?

This distribution-first strategy, and the "stationary OCRS" framework, are also useful for more general downward-closed feasibility environments [2]. For general graph matchings, the same maximum-entropy construction has an especially simple analysis: an edge can fail to be addable only if one of its endpoints is already matched, so a union bound gives a $1/3$ guarantee. An identical argument extends to rank-L hypergraph matchings, giving a $1/(L+1)$ guarantee (see also [3, 12]).

Our paper [2] also applies the same idea to rank-k uniform matroids, where feasibility means selecting at most k active elements, obtaining an asymptotically optimal selectability of $1 - \sqrt{2/(\pi k)} + O(1/k)$, significantly simplifying previous algorithms and analyses [1, 4]. The stationary-distribution viewpoint can turn the design of online contention resolution schemes into a search for the right distribution over feasible sets, often yielding algorithms that are both explicit and easy to reason about.

A Maximum entropy and the Gibbs distribution

We recall why the maximum-entropy distribution used in the algorithm has the Gibbs form (see also [14]). Let $p \in \mathrm{relint}(\mathrm{conv}\{\mathbf{1}_M : M \in \mathcal{M}\})$, i.e., p is in the relative interior of the convex hull of the indicator vectors corresponding to matchings. In our setting, where $p_e = \alpha x_e$, $x_e > 0$, and x is a fractional matching, this assumption is satisfied. Consider the entropy maximization problem

$$\begin{aligned} \max_{\lambda \in \mathbb{R}^{\mathcal{M}}_{\geq 0}} \quad & -\sum_{M \in \mathcal{M}} \lambda_M \log \lambda_M \\ \text{s.t.} \quad & \sum_{M \in \mathcal{M}} \lambda_M = 1, \\ & \sum_{M \ni e} \lambda_M = p_e \qquad \forall e \in E. \end{aligned}$$

The assumption that p lies in the relative interior ensures that there is a feasible point with full support; hence the optimizer of the entropy also has full support on $\mathcal{M}$, and so the KKT conditions apply. Let β be the multiplier for the normalization constraint and let θ_e be the multiplier for the marginal constraint of edge e. For every matching M, the first-order condition gives

$$-\log \lambda_M - 1 + \beta + \sum_{e \in M} \theta_e = 0.$$

Thus

$$\lambda_M = \exp(\beta - 1) \prod_{e \in M} \exp(\theta_e).$$

Writing $w_e := \exp(\theta_e) > 0$ and normalizing, we obtain

$$\lambda_M = \frac{1}{Z(w)} \prod_{e \in M} w_e, \qquad Z(w) = \sum_{N \in \mathcal{M}} \prod_{f \in N} w_f.$$

Hence the maximum-entropy distribution is a Gibbs distribution. Note that θ minimizes the convex dual

$$h(\theta) := \log Z(\theta) - \langle \theta, p \rangle.$$

Moreover,

$$(\nabla h(\theta))_e = \Pr_{M \sim \mu_\theta}[e \in M] - p_e.$$

Thus, the ability to compute $\Pr_{M \sim \mu_\theta}[e \in M]$ would let us run a standard convex optimization method to find θ approximately. This is formalized by the framework of [14] and [15].

For general matchings (and hence bipartite matchings) classical MCMC methods give an FPRAS for the matching partition function Z (see chapter 12 of [9]), and hence also the corresponding one-edge marginal $\Pr_{M \sim \mu_\theta}[e \in M]$. The running time is polynomial in the problem size and $1/\eta$ for relative error η. Plugging this into the results of [14] and [15] yields a randomized polynomial-time procedure for computing the Gibbs parameters θ up to accuracy ε, with overall running time polynomial in $1/\varepsilon$ and $\log(1/\delta)$ where δ is the allowed probability of error.

Once we have access to the Gibbs parameters, there is a rapidly mixing Markov chain whose stationary distribution is the Gibbs distribution; once again see Chapter 12 in [9]. Thus we can sample from the maximum-entropy distribution up to total-variation error ε in time polynomial in the instance size and $\log(1/\varepsilon)$.

B Proof of positive correlation of unmatched vertices

We prove the positive correlation statement in Proposition 2. Fix positive edge weights w. For any subgraph H of G, let

$$Z(H) := \sum_{M \in \mathcal{M}(H)} \prod_{e \in M} w_e$$

be its matching partition function. For a vertex z, let $G - z$ be the graph obtained by deleting z and all incident edges. Under the Gibbs distribution on G,

$$\mathbb{P}[z \text{ is unmatched}] = \frac{Z(G-z)}{Z(G)}$$

and

$$\mathbb{P}[u, v \text{ are both unmatched}] = \frac{Z(G-u-v)}{Z(G)}.$$

Thus it suffices to prove

$$Z(G)Z(G-u-v) \geq Z(G-u)Z(G-v).$$

We prove this by constructing a weight-preserving injection

$$\Phi : \mathcal{M}(G-u) \times \mathcal{M}(G-v) \hookrightarrow \mathcal{M}(G) \times \mathcal{M}(G-u-v).$$

Take a pair $(M_u, M_v) \in \mathcal{M}(G-u) \times \mathcal{M}(G-v)$ and consider the symmetric difference

$$H := M_u \triangle M_v.$$

Every component of H is an alternating path or cycle. Let C_v be the component containing v. Since M_v avoids v, any edge of C_v incident to v belongs to M_u. It follows that C_v cannot contain u. Indeed, if an alternating path from $v \in V$ to $u \in U$ existed, it would have odd length, and since it starts with an M_u-edge, it would also end with an M_u-edge incident to u, contradicting $M_u \in \mathcal{M}(G-u)$.

Now swap the two matchings on C_v:

$$M := (M_v \setminus C_v) \cup (M_u \cap C_v), \qquad N := (M_u \setminus C_v) \cup (M_v \cap C_v).$$

Then M is a matching of G. The matching N avoids u because C_v does not contain u and M_u avoids u. It also avoids v because, on C_v, it uses the M_v-edges, and M_v avoids v. Thus

$$(M, N) \in \mathcal{M}(G) \times \mathcal{M}(G-u-v).$$

The product of weights is preserved:

$$\prod_{e \in M_u} w_e \prod_{e \in M_v} w_e = \prod_{e \in M} w_e \prod_{e \in N} w_e.$$

The map is injective because, from (M, N), the component containing v in $M \triangle N$ identifies exactly where the swap occurred and swapping back recovers (M_u, M_v). Summing the preserved weights over all pairs in $\mathcal{M}(G-u) \times \mathcal{M}(G-v)$ gives

$$Z(G-u)Z(G-v) \leq Z(G)Z(G-u-v),$$

as needed.

References

[1] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

[2] Mohammad Reza Aminian, Rad Niazadeh, and Pranav Nuti. Stationary online contention resolution schemes, 2026. URL https://arxiv.org/abs/2603.21532.

[3] José Correa, Andrés Cristi, Andrés Fielbaum, Tristan Pollner, and S Matthew Weinberg. Optimal item pricing in online combinatorial auctions. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 126–139. Springer, 2022.

[4] Atanas Dinev and S Matthew Weinberg. Simple and optimal online contention resolution schemes for k-uniform matroids. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, page 39:1–39:23. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.

[5] Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Prophet matching with general arrivals. *Math. Oper. Res*, 47(2):878–898, 2022.

[6] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes with applications to bayesian selection problems. *SIAM Journal on Computing*, 50(2):255–300, 2021.

[7] Nick Gravin and Hongao Wang. Prophet inequality for bipartite matching: Merits of being simple and nonadaptive. *Mathematics of Operations Research*, 48(1):38–52, 2023.

[8] Ole J Heilmann and Elliott H Lieb. Theory of monomer-dimer systems. *Communications in mathematical Physics*, 25(3):190–232, 1972.

[9] Mark Jerrum and Alistair Sinclair. The markov chain monte carlo method: an approach to approximate counting and integration. In Dorit S. Hochbaum, editor, *Approximation Algorithms for NP-hard Problems*. PWS Publishing, 1996.

[10] Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms (ESA 2018)*, volume 112, page 57:1–57:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.

[11] Will Ma, Calum MacRury, and Pranav Nuti. Online matching and contention resolution for edge arrivals with vanishing probabilities. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 159–159, 2024.

[12] Will Ma, Calum MacRury, and Jingwei Zhang. Online contention resolution schemes for network revenue management and combinatorial auctions. In *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, pages 100:1–100:1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2026.

[13] Calum MacRury, Will Ma, and Nathaniel Grammel. On (random-order) online contention resolution schemes for the matching polytope of (bipartite) graphs. *Operations Research*, 73(2):689–703, 2025.

[14] Mohit Singh and Nisheeth K Vishnoi. Entropy, optimization and counting. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 50–59, 2014.

[15] Damian Straszak and Nisheeth K Vishnoi. Maximum entropy distributions: Bit complexity and stability. In *Conference on Learning Theory*, pages 2861–2891. PMLR, 2019.