An Economic-Based Analysis of RANKING for Online Bipartite Matching

Alon Eden Michal Feldman Amos Fiat Kineret Segal

Abstract

We give a simple proof showing that the RANKING algorithm introduced by Karp, Vazirani and Vazirani [1] is $1 - \frac{1}{e}$ competitive for the online bipartite matching problem. Our proof resembles the proof given by Devanur, Jain and Kleinberg [2], but does not make an explicit use of linear programming duality; instead, it is based on an economic interpretation of the matching problem. In our interpretation, one set of vertices represent items that are assigned prices, and the other set of vertices represent unit-demand buyers that arrive sequentially and choose their most-demanded items.

1 Problem Statement

Consider a bipartite graph $G = (L, R; E)$, where $L = \{\ell_1, \ldots, \ell_n\}$ and $R = \{r_1, \ldots, r_n\}$ are the left and right vertices, respectively, and E is the set of edges.

The online bipartite matching problem introduced by Karp, Vazirani and Vazirani [1] is the following: The graph G is initially unknown. In iteration i, for $i = 1, \ldots, n$, vertex ℓ_i arrives, along with its adjacent edges (which are unknown from the outset). The algorithm needs to decide which neighbor of ℓ_i (if any) ℓ_i is matched to; this decision is irrevocable. The objective is to maximize the cardinality of the obtained matching.

A simple greedy algorithm for this problem matches each arriving vertex with an arbitrary unmatched neighbour, if available. Every greedy algorithm outputs a maximal matching, hence has cardinality at least a half of the maximum matching. It is not very difficult to see that this bound is tight; that is, there exist graphs for which this greedy algorithm cannot achieve more than half of the maximum matching. It was also shown that a randomized version of the greedy algorithm, which chooses a currently unmatched neighbor uniformly at random (if one exists) also has a competitive ratio of $1/2$, up to low order terms [1].

In [1], Karp et al. introduced the randomized RANKING algorithm, and proved that it has a competitive ratio of $1 - \frac{1}{e}$. They also showed that this bound is tight (up to low order terms). RANKING first chooses a random permutation π over the vertices in R. Upon the arrival of a vertex

ℓ_i, RANKING matches ℓ_i to the highest-ranked (according to π) currently unmatched neighbor of ℓ_i.

The analysis in the original paper was quite complicated. Subsequent papers by Goel and Mehta [3], Birnbaum and Mathieu [4] and Devanur Jain and Kleinberg [2] simplified the analysis considerably. The proof presented here is based on an economic interpretation of the online bipartite matching problem. It is similar to the proof of [2], but does not make an explicit use of linear programming duality.

2 An Economic-Based Analysis of RANKING

Consider the following interpretation of the RANKING algorithm. Given a graph $G = (L, R; E)$, vertices of R represent items, and vertices of L represent utility maximizing buyers. If $(\ell_i, r_j) \in E$, then we say that buyer ℓ_i is connected to item r_j. Every buyer ℓ_i has a binary unit-demand valuation, with value 1 to items connected to ℓ_i and value 0 otherwise.

Before the arrival of buyers, every item r_j is assigned a price $p_j = e^{w_j - 1}$, where w_j is a uniformly random number in $[0, 1]$ (chosen independently for every item). Buyers arrive in arbitrary order. Every buyer, upon arrival, chooses an item that maximizes her *utility*, defined as the difference between her value for the item and the item's price. This means that every buyer ℓ_i chooses the cheapest item she is connected to, which is still available.

We claim that the market process above is equivalent to the RANKING algorithm. To see this, one needs to show that every buyer purchases the item that is ranked highest among all available items, according to a preset random permutation. In the market setting, every buyer purchases the cheapest (currently available) item she is connected to. But since the price of every item is a strictly monotonically increasing function of w_j, which is chosen independently and uniformly at random, the permutation induced by item prices is a random permutation.

We now proceed to the analysis of the market process.

For each item r_j, let rev_j denote the revenue obtained by r_j (i.e., p_j if the item was purchased and 0 otherwise). The utility of buyer ℓ_i is

$$\text{util}_i = \begin{cases} 1 - p_j, & \text{if buyer } \ell_i \text{ purchased item } r_j \\ 0, & \text{if buyer } \ell_i \text{ did not purchase any item} \end{cases}$$

Fix some arrival order of the buyers and a price vector $\mathbf{p} = (p_1, \ldots, p_n)$, and let T be the set of the corresponding purchased items. Since every buyer that received an item has value 1, the social welfare is the cardinality of the obtained matching. The following equation shows that it can also be written as the sum of the buyers' utilities and the total revenue:

$$\sum_{\ell_i \in L} \text{util}_i + \sum_{r_j \in R} \text{rev}_j = \sum_{r_j \in T} (1 - p_j) + \sum_{r_j \in T} p_j = \sum_{r_j \in T} 1 = |T|. \tag{1}$$

We note that the approach of expressing the welfare as the sum of utilities and revenue has been used previously in other settings and proved useful [5, 6, 7].

Recall that for weights $\mathbf{w} := (w_1, \ldots, w_n)$ we set prices $p_j = e^{w_j - 1}$. We shall now present the main claim of the proof.

Claim 2.1. *For every order of arrival of the buyers, let w_j be chosen uniformly in $[0, 1]$, and let prices be as above. We have that for all edges $(\ell_i, r_j) \in E$:*

$$\mathop{\mathbb{E}}_{\mathbf{w}} [\text{util}_i + \text{rev}_j] \geq 1 - \frac{1}{e}.$$

Before proving Claim 2.1, we show how it is used to prove the competitive ratio of $1 - \frac{1}{e}$. Fix a maximum matching M^* and let M be the matching produced by the market process above. It follows that

$$\begin{aligned} \mathop{\mathbb{E}}_{\mathbf{w}} [|M|] &= \mathop{\mathbb{E}}_{\mathbf{w}} \left[\sum_i \text{util}_i + \sum_j \text{rev}_j \right] \geq \mathop{\mathbb{E}}_{\mathbf{w}} \left[\sum_{(\ell_i, r_j) \in M^*} \text{util}_i + \text{rev}_j \right] \\ &= \sum_{(\ell_i, r_j) \in M^*} \mathop{\mathbb{E}}_{\mathbf{w}} [\text{util}_i + \text{rev}_j] \geq \left(1 - \frac{1}{e}\right) |M^*|, \end{aligned}$$

where the first equality follows from Equation (1), and the last inequality follows from Claim 2.1.

We now proceed to proving Claim 2.1 — thus proving the competitive ratio of RANKING.

Proof of Claim 2.1. Fix some arbitrary order of buyer arrival σ, buyer ℓ_i item r_j such that $(\ell_i, r_j) \in E$, and let prices $\mathbf{p}$ be random prices as above. Consider the market without item r_j and let $p = e^{y-1}$ be the price of the item chosen by buyer ℓ_i under the same arrival order σ (if ℓ_i buys nothing, set $p = 1$). Then, under order σ, in the market with item r_j, we have the two following properties:

1. Item r_j is always sold when $p_j < p$. This follows since either (a) some previous buyer bought item r_j, or (b) buyer ℓ_i prefers item r_j over the item chosen by ℓ_i when r_j was unavailable.

2. The utility of buyer ℓ_i, $\text{util}_i \geq 1 - p$. Observe that after reintroducing item r_j to the market, every buyer has the same set of items available to him plus — possibly — one additional item. This is obviously true for the first incoming buyer, and remains true subsequently since the introduction of an additional item does never induces a buyer to take an item previously waived.

Property (1) above implies that

$$\mathop{\mathbb{E}}_{\mathbf{w}}\left[\mathrm{rev}_j\right] = \mathop{\mathbb{E}}_{\mathbf{w}}\left[p_j \cdot \mathbb{1}_{r_j \text{ is sold}}\right] \geq \mathop{\mathbb{E}}_{\mathbf{w}}\left[p_j \cdot \mathbb{1}_{p_j<p}\right] = \int_0^y e^{w_j-1} dw_j = e^{y-1} - \frac{1}{e} = p - \frac{1}{e}.$$

It now follows from property (2) that

$$\mathop{\mathbb{E}}_{\mathbf{w}}\left[\mathrm{util}_i + \mathrm{rev}_j\right] \geq 1 - p + p - \frac{1}{e} = 1 - \frac{1}{e},$$

as desired.

□

Remark 1. Devanur et al. [2] gave an elegant and simple proof of the $1 - 1/e$ bound achieved by RANKING using primal-dual analysis, where the primal LP represents the matching problem. It is known that the dual variables can be interpreted as these util_i's and rev_j's. Our proof uses a scaled version of the assignment of the relevant dual variables in [2] as prices for items. The new interpretation simplifies the proof in two ways. (a) It removes the need to argue about the dual program and its feasibility. (b) Viewed from the economic perspective, some of the arguments in [2] are more readily apparent.

Remark 2. Note that the proof Claim 2.1 only uses the random choice of w_j, while all other values can be arbitrary.

Remark 3. While the choices of the buyers, and thus the $1 - \frac{1}{e}$ bound, hold when prices are just chosen uniformly at random in $[0, 1]$[1], the proof of Claim 2.1 requires that we use the prices as specified above. Specifically, consider the lower bound example from [1]. In this instance, the last buyer to arrive is very unlikely to receive anything (in particular, it must be the case that the price of the last item is maximal, which happens w.p. $1/n$). Claim 2.1 is about all edges, in particular the edge from the last buyer. As the utility of the last buyer is small, the revenue must compensate, but the revenue from this item is at most $\sim 1/2$ under uniform price distributions.

Acknowledgements. This work was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122, and by the Israel Science Foundation (grant number 317/17).

References

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

[1] or an arbitrary bijection from $\mathbf{w}$ to $[0, 1]$

[2] Nikhil R. Devanur, Kamal Jain, and Robert D. Kleinberg. Randomized primal-dual analysis of RANKING for online bipartite matching. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 101–107, 2013.

[3] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adwords. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, pages 982–991, 2008.

[4] Benjamin E. Birnbaum and Claire Mathieu. On-line bipartite matching made simple. *SIGACT News*, 39(1):80–87, 2008.

[5] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 123–135, 2015.

[6] Paul Duetting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 540–551, 2017.

[7] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 700–714, 2018.