DS 574 Algorithmic Mechanism Design	**Lecture #14 Worksheet**
Fall 2023	Prof. Kira Goldner

Online Bipartite Matching [KVV '90, EFFS '21]

In the online bipartite matching setting [Karp, Vazirani, and Vazirani, 1990], there is a bipartite graph $G = (L \cup R, E)$ where are vertices are split into the left side, L, and the right side, R. Edges are unweighted, i.e., all have a weight of 1. We are in an online setting where we see R up front, but the vertices of L arrive online, and as each vertex arrives, we see which edges are incident to it from R. The objective is to match vertices in L to those in R, immediately and irrevocably as each vertex arrives, forming a matching M, such that we maximize the cardinality of the matching $|M|$ and compare well to the maximal offline matching.

In the original paper, Karp et al. [1990] show:

- For every deterministic algorithm, $|M| \leq n/2$.
- Choosing a random match for each vertex independently implies that $\mathbb{E}[|M|] \leq n/2$.
- RANKING (KVV): Choosing a global ranking π U.A.R. and matching according to π implies that $\mathbb{E}[|M|] \geq (1 - 1/e)n$.
- This is tight!

The first proof [Karp et al., 1990] was very complicated (and imprecise). Simplifications were given by Goel and Mehta [2008], Birnbaum and Mathieu [2008], Devanur, Jain, and Kleinberg [2013].

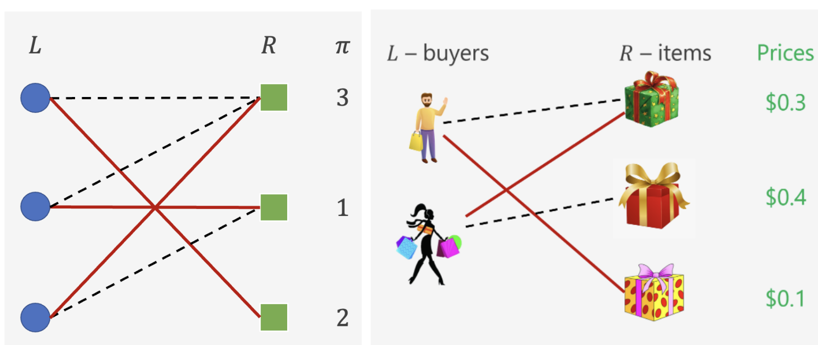

Figure 1: Online Bipartite Matching: RANKING and its economic interpretation.

Today, we'll look at an equivalent algorithm and simplified analysis due to Eden, Feldman, Fiat, and Segal [2021]. We'll use the following set-up: Let R be items and L buyers. They have value 1 or 0 for each item (depending on whether there is an edge). They are unit-demand (want one match). Then our algorithm is as follows.

-

-

Observations:

- For any F supported on $(0,1)$ without pointmasses, choosing item prices i.i.d. from F is equivalent to:

- The *welfare* of the matching is:

  We can rewrite it as:

Lemma 1. *For F that samples $w \sim U[0,1]$ and sets $p_j = e^{w-1}$, we have for every buyer i and item j such that (i,j) is an edge in M:*

$$\mathbb{E}[util_i + rev_j] \geq 1 - 1/e.$$

Corollary 1. $|M| \geq$

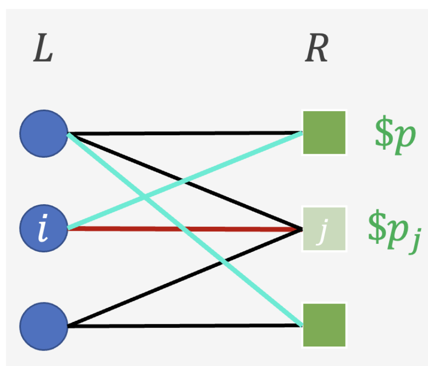

Figure 2: Economic analysis of RANKING.

Proof.

Conclusions:

- Taking the economic perspective can be very useful in algorithm analysis.
- Decomposing value into revenue + utility is a powerful tool, e.g., [Feldman, Gravin, and Lucier, 2015, Dütting, Feldman, Kesselheim, and Lucier, 2017, Ehsani, Hajiaghayi, Kesselheim, and Singla, 2018].

Robustness: Prior-Independence

"Prior-independent" results give us guarantees in the event that the designer *doesn't* know the distribution F from which the bidders' values are drawn. In this case, we assume that their values are still drawn from a prior distribution, as in the Bayesian setting, so there is some revenue-optimal mechanism $\textsc{opt}(F)$ that we wish to approximate, we just have to do so without knowing F.

The Bulow-Klemperer Result

One famous result takes the form of resource augmentation.

Theorem 2 (Bulow and Klemperer [1994])**.** *For i.i.d. regular single-item environments, the expected revenue of the second-price auction with $n+1$ agents is at least that of the optimal auction with n agents.*

Let's talk about what this theorem is saying. Instead of finding the optimal auction tailored to a distribution F for n agents, you can use the Vickrey auction, which requires no prior knowledge of the distribution, so long as:

-

-

This result **does not hold** without these assumptions. However, it is a *very strong* result, should our setting meet these assumptions.

Proof.

The Single Sample Mechanism

Can't recruit extra buyers? Instead, we can just exclude one. This is what the single sample result says.

Theorem 3 (Dhangwatnotai, Roughgarden, and Yan [2015])**.** *Given a random sample from a bidder's distribution, posting it as a take-it-or-leave-it price gives a* $\frac{1}{2}$*-approximation to the optimal revenue.*

Figure 3: Geometric intuition for a posted-price from a single sample.

Proof. In quantile space!

It turns out, using a single sample from the buyers' distribution to set reserve prices and running VCG is a good approximation to the optimal mechanism. See Hartline chapter 5 for more.

References

Benjamin Birnbaum and Claire Mathieu. On-line bipartite matching made simple. *Acm Sigact News*, 39(1):80–87, 2008.

Jeremy I Bulow and Paul D Klemperer. Auctions vs. negotiations, 1994.

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*, pages 101–107. SIAM, 2013.

Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 540–551. IEEE Computer Society, 2017. doi: 10.1109/FOCS.2017.56. URL `https://doi.org/10.1109/FOCS.2017.56`.

Alon Eden, Michal Feldman, Amos Fiat, and Kineret Segal. An economics-based analysis of ranking for online bipartite matching? In *Symposium on Simplicity in Algorithms (SOSA)*, pages 107–110. SIAM, 2021.

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*, pages 700–714. SIAM, 2018.

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*, pages 123–135. SIAM, 2015.

Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adwords. In *SODA*, volume 8, pages 982–991, 2008.

Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 352–358, 1990.