Optimal Algorithm for Max-Min Fairness Bandit

Zilong Wang[1], Zhiyao Zhang[2], **Shuai Li**[1]

[1] Shanghai Jiao Tong University

[2] Ohio State University

Max-Min Fairness for Multiple Players

[Radunovic et al. 2007, Asadpour et al. 2010, Zahavi et al. 2013]

☺	1	0.8	0.8	0.6
☺	1	0.9	0.8	0.5
☹	0.8	0.8	0.6	0.3

- maximize total reward maybe unfair to certain player

☺	1	0.8	0.8	0.6
☺	1	0.9	0.8	0.5
☺	0.8	0.8	0.6	0.3

- max-min fairness maximizes the reward of the minimal player

- Max-min matching $m^* \in \arg\max_m \min_i \mu_{i,m_i}$ with

  Max-min value $\gamma^* \;=\; \max_m \min_i \mu_{i,m_i}$ utilities

- Usually assume #players ≤ #items

(Offline) Algorithm: Known Utilities $\{\mu_{i,k}\}$

- N players, K items, $N \le K$
- Binary Search Threshold Test Algorithm [Panagiotas et al. 2023]
    - Sort $\{\mu_{i,k}\}_{i\in[N],k\in[K]}$
    - Binary search candidate max-min value γ:
        - Delete edges (i,k) with $\mu_{i,k} < \gamma$
        - Run Hungarian algorithm for max matching
        - If max matching has size N, set γ as lower bound, continue search

	1	0.8	0.8	~~0.6~~
	1	0.9	0.8	~~0.5~~
	0.8	0.8	~~0.6~~	~~0.3~~

- Sort $\mu_{i,k}$: 1, 0.9, 0.8, 0.6 ,0.5, 0.3
- Test $\gamma = 0.8$
    - Delete edges with $\mu_{i,k} < \gamma = 0.8$
    - Find max-matching of size $= 3 = N$ Yes!

(Offline) Algorithm: Known Utilities $\{\mu_{i,k}\}$ 2

- N players, K items, $N \leq K$
- Binary Search Threshold Test Algorithm [Panagiotas et al. 2023]
  - Sort $\{\mu_{i,k}\}_{i \in [N], k \in [K]}$
  - Binary search candidate max-min value γ:
    - Delete edges (i, k) with $\mu_{i,k} < \gamma$
    - Run Hungarian algorithm for max matching
    - If max matching has size N, set γ as lower bound, continue search
    - Else set γ as upper bound, continue search
- Time complexity: $O(\log(NK)\, NK^2)$

The complexity of Hungarian algorithm is $O((N+K)NK)$. [Kuhn 1955]

1	~~0.8~~	~~0.8~~	~~0.6~~
1	0.9	~~0.8~~	~~0.5~~
~~0.8~~	~~0.8~~	~~0.6~~	~~0.3~~

- Sort $\mu_{i,k}$: 1, 0.9, 0.8, 0.6 ,0.5, 0.3
- Test $\gamma = 0.8$
  - Delete edges with $\mu_{i,k} < \gamma = 0.8$
  - Find max-matching of size $= 3 = N$ Yes!
- Test $\gamma = 0.9$
  - Delete edges with $\mu_{i,k} < \gamma = 0.9$
  - Find max-matching of size $= 2 < N$ No!

$\Rightarrow \gamma^* = 0.8$

Multi-Player Multi-Armed Bandit

- N players, K arms, $N \le K$
- At round $t = 1,2,\dots,T$:
    - Each player $i \in [N]$ selects an arm $m_i(t) \in [K]$ and observes
        - No collision: 1-subgaussian reward $X_i(t)$ with unknown mean $\mu_{i,m_i(t)}$
        - Collision: 0 reward
    - Player i cannot observe other players' actions and rewards
    - (Decentralized) Each player updates him/herself

- Objective: Minimize max-min regret:

$$R(T) = \mathbb{E}\left[T\gamma^* - \sum_{t=1}^{T} \min_{i \in [N]} X_i(t)\right]$$

Existing Results

	Regret Type	Regret Bound
Bistritz et al. ICML 2020	Upper Bound	$O(\exp(NK)\log T \log\log T + \exp\frac{1}{\Delta_1})$
	Lower Bound	$\Omega\left(\frac{\log T}{\Delta}\right)$
Leshem. ICASSP 2025	Upper Bound	$O(N^2K\log T \log\log T + \exp\frac{1}{\Delta_1})$

huge gap!

$\Delta_1 := \min_{i,i',k,k'} |\mu_{i,k} - \mu_{i',k'}|$: minimum gap among all (i,k) pairs (require all pairs are distinct)

$\Delta := \min_{(i,k):\mu_{i,k}<\gamma^*} (\gamma^* - \mu_{i,k})$: minimum gap between max-min value and sub-optimal pair

Problem Hardness? Lower Bound

- Each sub-optimal pair (i, k) needs to be explored sufficient $\Omega\left(\frac{\log T}{\Delta^2}\right)$ times
- But each round can explore multiple ($\leq N$) pairs with a matching
- There are NK pairs in total
- What is the problem hardness for the coefficient in front of $\Omega\left(\frac{\log T}{\Delta^2}\right)$?
  - NK or N^2 or K ?
- Can we design hard instance?

Lower Bound - Instance Design

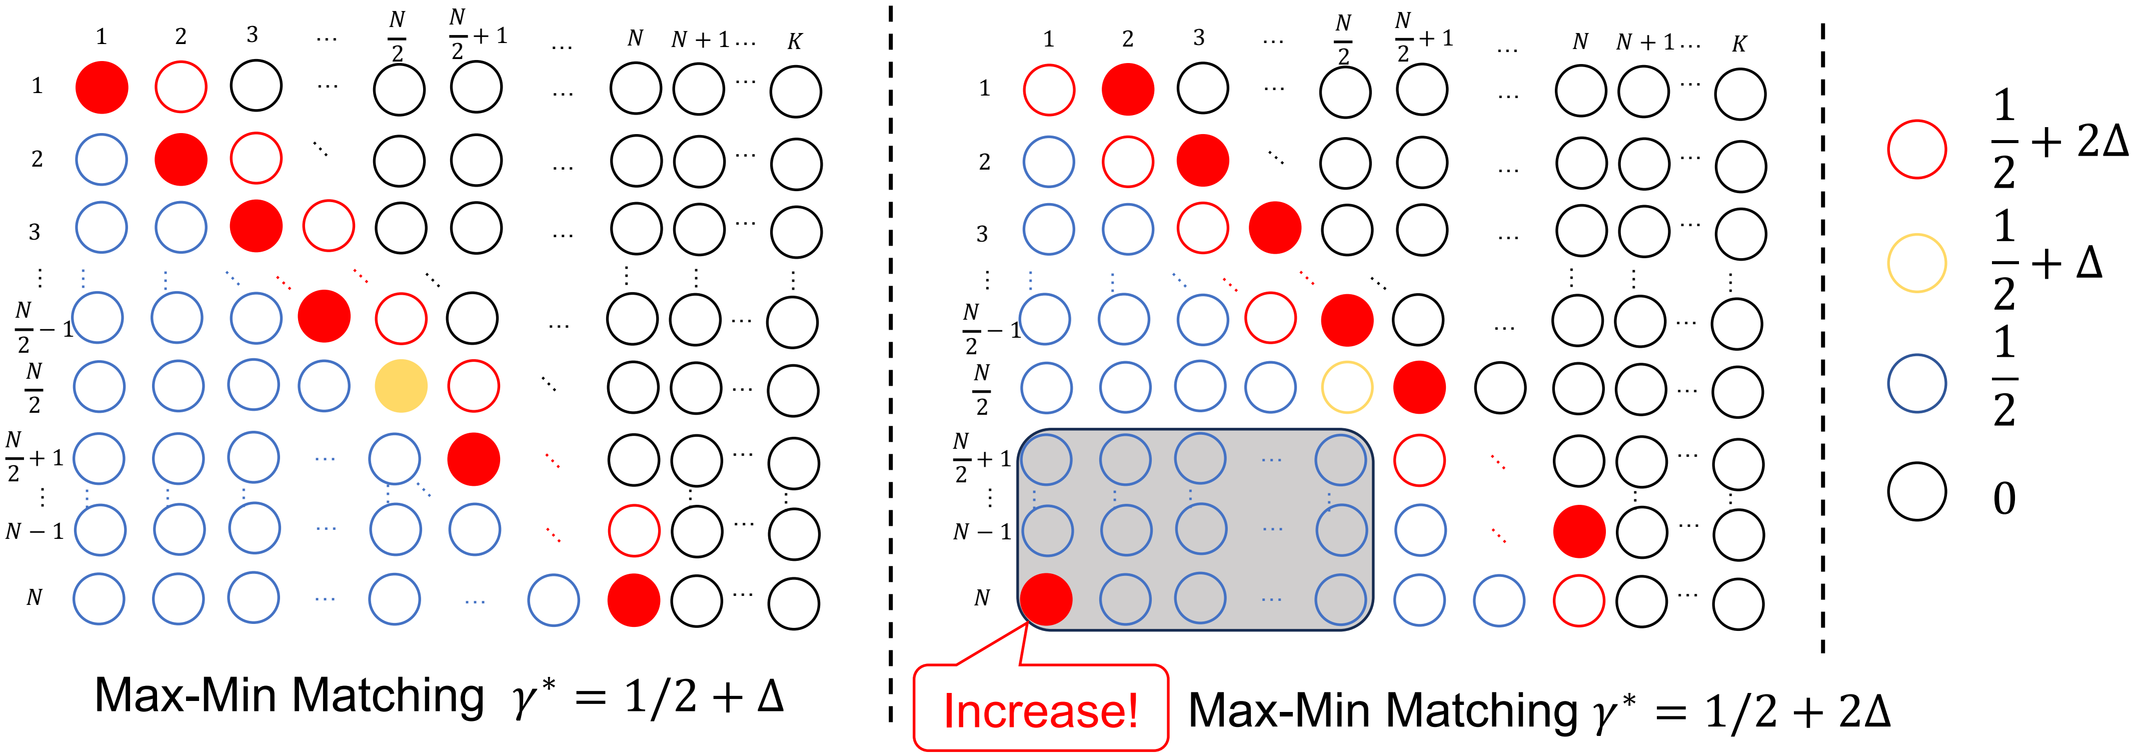

Max-Min Matching $\gamma^* = 1/2 + \Delta$

Max-Min Matching $\gamma^* = 1/2 + 2\Delta$

- Increasing any pair in Gray-Box would change m^* and γ^*
- $\Longrightarrow$ Each pair needs to explore $N_{i,k} \geq \Omega\left(\frac{\log T}{\Delta^2}\right)$
- Key Observation: $\exists O(N^2)$ pairs in Gray-Box , use this to prove coefficient of $O(N^2)$

Lower Bound Proof

- $\text{Regret} = \mathbb{E}\left[T\gamma^* - \sum_{t=1}^{T} \min_{i \in [N]} X_i(t)\right]$

$$= \Delta \cdot \mathbb{E}\left[\underbrace{\sum_{t=1}^{T} \mathbb{I}\left\{\min_i \mu_{i,m_i(t)} = 1/2\right\}}_{N_1}\right] + \gamma^* \cdot \mathbb{E}\left[\underbrace{\sum_{t=1}^{T} \mathbb{I}\left\{\min_i \mu_{i,m_i(t)} = 0\right\}}_{N_2}\right]$$

suboptimal matching without selecting 0-pair

- Claim: Without selecting 0-pair, at most one pair in Gray-Box will be selected

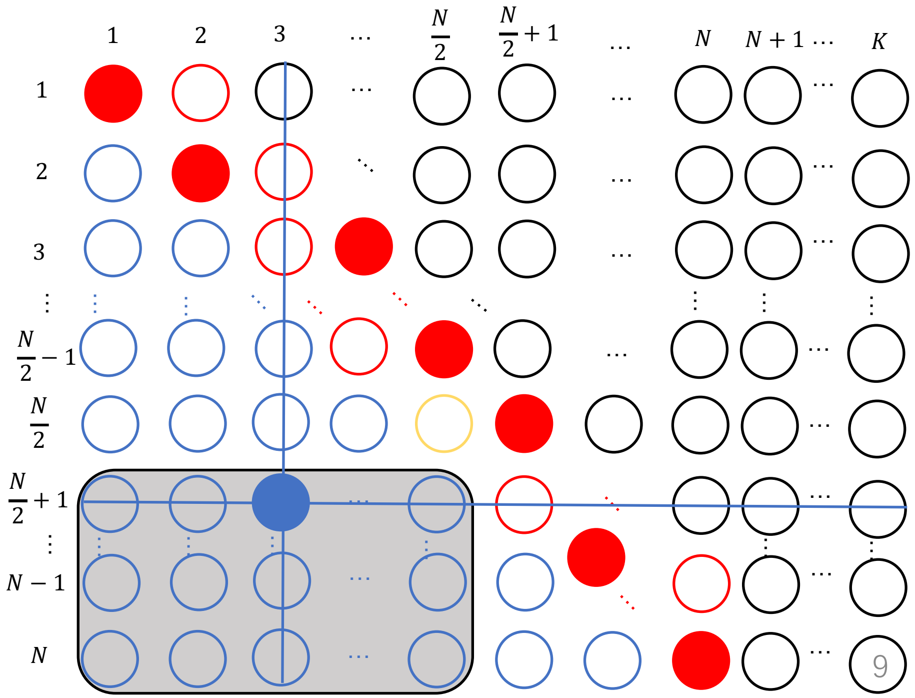

- N_1 = #{one 1/2-pair is selected and no 0-pair}
- N_2 = #{0-pair is selected} ⇒ at most N ½-pair is selected
- $\Omega\left(\frac{N^2 \log T}{\Delta^2}\right) = \sum_{i \in (N/2, N]} \sum_{k \in [1, N/2]} N_{i,k} \le N_1 + N \cdot N_2$
  $\le N_1 + \frac{\gamma^*}{\Delta} \cdot N_2$ = Regret / Δ
- $\Longrightarrow$ Regret= $\Omega\left(\frac{N^2 \log T}{\Delta}\right)$

$\Delta \le \gamma^*/N$

Lower Bound Proof 2

- $\forall$ pair (N, k) with $k > N$, increasing its value changes m^* and γ^*
- Each such pair needs to explore $\Omega\left(\frac{\log T}{\Delta^2}\right)$ times
- Each round explores at most one such pair $\Longrightarrow \Omega\left(\frac{(K-N)\log T}{\Delta}\right)$ regret

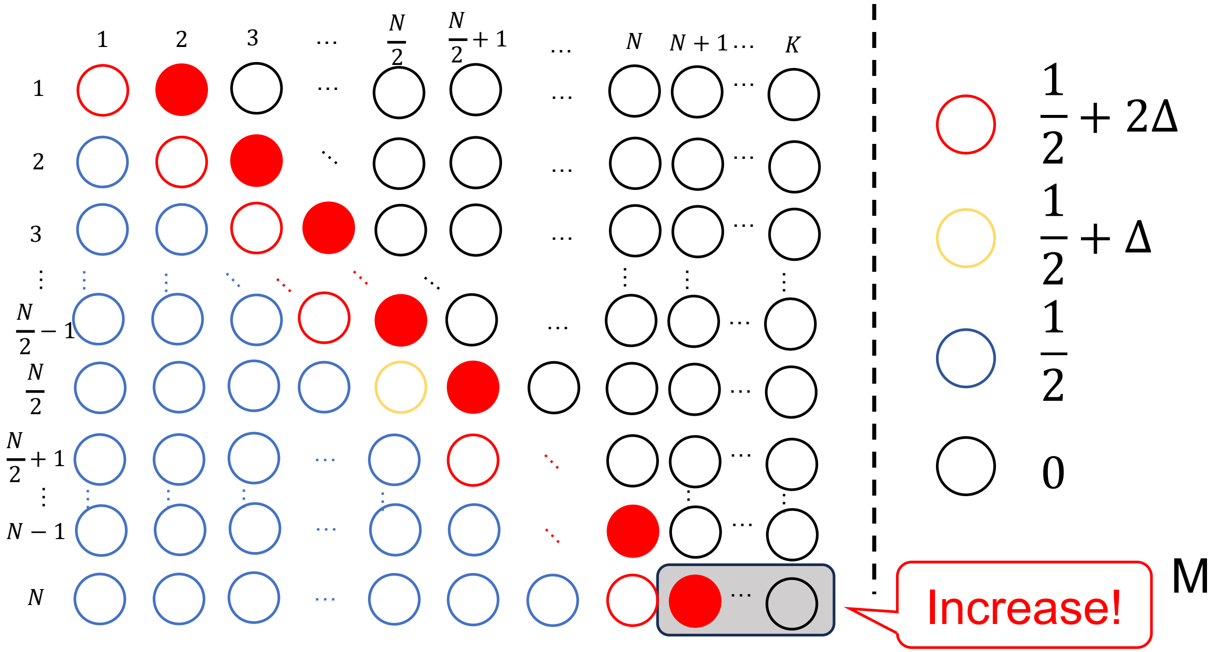

Max-Min Matching
$\gamma^* = 1/2 + 2\Delta$

Lower Bound

- **Theorem (Lower Bound)** N players, K arms, time horizon T, $\exists$ an instance with $\Delta \leq \gamma^*/N$, s.t. any uniformly consistent algorithm satisfies the max-min regret

$$R(T) \geq \Omega\left(\frac{\max\{N^2, K\} \log T}{\Delta}\right)$$

Online Algorithm: Decentralized Fair Elimination

- Phased-based exploration
- At each phase s:
  - **Exploration** s.t. each remaining pair is selected for 2^s times
  - Players communicate their statistics: $\{\hat{\mu}_{i,k}, N_{i,k}\}_{i,k}$
  - **Eliminate** unnecessary pairs

Online Algorithm: Exploration

- Begin w/ any candidate matching m using remaining pairs
  - Such a matching must exist if only sub-optimal pairs are deleted
- Denote arms in m as candidate arms

Use any arm to make up a full matching

$\Longrightarrow N^2$ rounds to explore each pair once

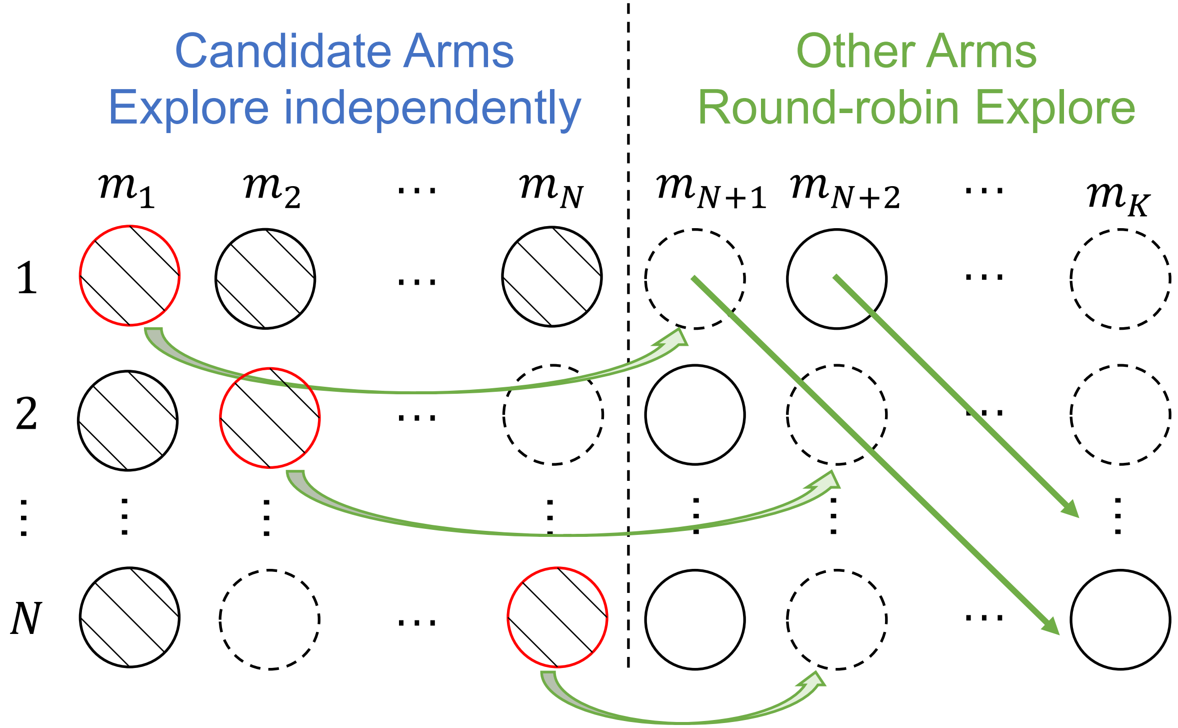

Use candidate match to make up a full matching

$\Longrightarrow K$ rounds to explore other arms

$\Longrightarrow N^2 + K$ rounds to explore each remaining pair once
$\Longrightarrow (N^2 + K)2^s$ rounds in phase s

Online Algorithm: Elimination

- For (i,k), $\mathrm{UCB}_{i,k}(t) = \hat{\mu}_{i,k}(t) + c\sqrt{\frac{\log T}{N_{i,k}(t)}}$, $\mathrm{LCB}_{i,k}(t) = \hat{\mu}_{i,k}(t) - c\sqrt{\frac{\log T}{N_{i,k}(t)}}$
- Compute max-min value γ_{LCB} using $\mathrm{LCB}_{i,k}$
  - $\Longrightarrow \gamma_{\mathrm{LCB}} \leq \gamma^*$ w.h.p.
- Eliminate pairs (i,k) if
  - $\mathrm{UCB}_{i,k} < \gamma_{\mathrm{LCB}}$, or
  - (i,k) cannot form a matching

(i,k) will be eliminated when $N_{i,k} \geq \frac{\log T}{(\Delta_{i,k}/2)^2}$

before phase s: $2^s \geq \frac{\log T}{(\Delta_{i,k}/2)^2}$

- Regret in phase $s \leq (N^2+K)2^s \cdot \max_{(i,k)\in\mathrm{Phase}\ s} \Delta_{i,k} \leq O\left((N^2+K)2^s \cdot \sqrt{\frac{\log T}{2^s}}\right)$
- Regret $\leq O\left((N^2+K)\sum_s \sqrt{2^s \log T}\right) \leq O\left((N^2+K)\frac{\log T}{\Delta}\right)$

$s_{\max} \leq O\left(\log\frac{\log T}{\Delta^2}\right)$

Online Algorithm: Regret

- **Theorem (Upper Bound)** N players, K arms, time horizon T. The max-min regret of the *Decentralized Fair Elimination* algorithm satisfies

$$R(T) \leq O\left(\frac{(N^2 + K)\log T}{\Delta} + C_{\text{comm}} \log T\right)$$

phase will go on if m^* is not unique

Result

	Regret Type	Regret Bound
Bistritz et al. ICML 2020	Upper Bound	$O(\exp(NK)\log T \log\log T + \exp\frac{1}{\Delta_1})$
	Lower Bound	$\Omega\left(\frac{\log T}{\Delta}\right)$
Leshem. ICASSP 2025	Upper Bound	$O(N^2K\log T \log\log T + \exp\frac{1}{\Delta_1})$
Our Work	Upper / Lower Bound	$\Theta\left((N^2+K)\frac{\log T}{\Delta}\right)$

Close the gap!

$\Delta_1 := \min_{i,i',k,k'} |\mu_{i,k} - \mu_{i',k'}|$: minimum gap among all (i,k) pairs (require all pairs are distinct)

$\Delta := \min_{(i,k):\mu_{i,k}<\gamma^*} (\gamma^* - \mu_{i,k})$: minimum gap between max-min value and sub-optimal pair

Conclusion

- Lower bound: Design worst case instance
- Upper bound: Propose optimal exploration assignment strategy
- Future Direction:
  - Strategy-proof of players
  - Robustness against attacker
  - Extend the analysis to other notions of fairness, e.g., envy-freeness [Gamow et al. 1958] and proportionality [Steinhaus 1948]

Thanks!
&
Questions?

Shuai Li

- Associate professor at
  Shanghai Jiao Tong University
- Research interests: RL/ML Theory
- Personal website: https://shuaili8.github.io/

References

- [Asadpour et al. 2010] Asadpour, A. and Saberi, A. An approximation algorithm for max-min fair allocation of indivisible goods. SIAM Journal on Computing, 39(7):2970–2989, 2010.
- [Radunovic et al. 2007] Radunovic, B. and Le Boudec, J.-Y. A unified framework for max-min and min-max fairness with applications. IEEE/ACM Transactions on networking, 15(5): 1073–1083, 2007.
- [Zahavi et al. 2013] Zehavi, E., Leshem, A., Levanda, R., and Han, Z. Weighted max-min resource allocation for frequency selective channels. IEEE transactions on signal processing, 61 (15):3723–3732, 2013.
- [Panagiotas et al. 2023] Panagiotas, I., Pichon, G., Singh, S., and U c¸ ar, B. Engineering Fast Algorithms for the Bottleneck Matching Problem. In Gørtz, I. L., Farach-Colton, M., Puglisi, S. J., and Herman, G. (eds.), 31st Annual European Symposium on Algorithms (ESA 2023)
- [Bistritz et al. 2020] Bistritz, I., Baharav, T., Leshem, A., and Bambos, N. My fair bandit: Distributed learning of max-min fairness with multi-player bandits. In International Conference on Machine Learning, pp. 930–940. PMLR, 2020.
- [Leshem. 2025] Leshem, A. Near optimal privacy preserving fair multi-agent bandits*. In ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1–5, 2025. doi: 10.1109/ICASSP 49660. 2025.10889464.
- [Gamow et al. 1958] Gamow, George; Stern, Marvin (1958). Puzzle-math. Viking Press. ISBN 0670583359.
- [Steinhaus 1948] Steinhaus, Hugo (1948). "The problem of fair division". Econometrica. 16 (1): 101–104. JSTOR 1914289.