Proportionally Fair Makespan Approximation

Tomek Ponitka
Tel Aviv University

Joint work with

Michal Feldman
Tel Aviv University
Microsoft ILDC

Jugal Garg
University of Illinois
Urbana-Champaign

Vishnu V. Narayan
University of Illinois
Urbana-Champaign

Setting

Model: Scheduling n jobs on m unrelated machines

Objective: Minimizing makespan


	client 1	client 2
salesman 1	2 hours	3 hours
salesman 2	$1-\epsilon$ hours	2 hours

Makespan = 2 hours

Algorithmic Aspects

Scheduling is among the most extensively studied problems in computer science.

Can we approximate makespan in polynomial time?

Theorem [Lenstra, Shmoys, Tardos, FOCS 1990]:
There exists a 2-approximate polynomial-time algorithm.

Unless P = NP, no polynomial-time algorithm can achieve a better than 3/2-approximation.

Machine Utilities

Definition: A **mechanism** (X, t) is a job allocation to machines X together with **transfers** t among machines.

		client 1	client 2	transfer	total cost
	salesman 1	2 hours	3 hours	½ hour	2½ hours
	salesman 2	$1 - \epsilon$ hours	2 hours	−½ hour	1½ hours

Makespan = 2 hours

Scheduling and Incentives

Scheduling is a fundamental problem in algorithmic mechanism design.

How well can a strategy-proof mechanism approximate makespan?

Theorem [Nisan, Ronen, STOC 1999]:
There exists an m-approximate strategy-proof mechanism.

Theorem [Christodoulou, Koutsoupias, Kovács, STOC 2023]:
Any strategy-proof mechanism is at least m-approximate.

[Christodoulou, Koutsoupias, Vidali, SODA 2007] [Koutsoupias, Vidali, MFCS 2007]
[Dobzinski, Shaulker, 2020] [Christodoulou, Koutsoupias, Kovács, FOCS 2022]

Scheduling and Fairness

Our Main Focus:

Can we ensure fairness for machines?

Mechanisms with transfers resemble fair division with subsidy.

[Aragones, 1995] [Halpern, Shah, SAGT 2019] [Brustle, Dippel, Narayan, Suzuki, Vetta, EC 2020]
[Barman, Krishna, Narahari, Sadhukhan, IJCAI 2022] [Caragiannis, Ioannidis, WINE 2022]
[Goko, Igarashi, Kawase, Makino, Sumita, Tamura, Yokoi, Yokoo, AAMAS 2022]
[Choo, Ling, Suksompong, Teh, Zhang, ORL 2024] [Wu, Zhang, Zhou, WINE 2024]

How well can an envy-free mechanism approximate makespan?

Theorem [Cohen, Feldman, Fiat, Kaplan, Olonetsky, EC 2010]:
There exists an $O(\log m)$-approximate envy-free mechanism.

[Hartline, Ieong, Mu'alem, Schapira, Zohar, ADT 2009]

Envy-Freeness

Definition: A mechanism (X, t) is envy-free if

$c_i(X_i) + t_i \leq c_i(X_j) + t_j$ for any machines $i, j \in [m]$


	client 1	client 2	transfer	total cost
salesman 1	2 hours	3 hours	½ hour	2½ hours
salesman 2	$1 - \epsilon$ hours	2 hours	−½ hour	1½ hours

Salesman 2 envies salesman 1 since: 2 – ½ > $(1 - \epsilon)$ + ½.

Envy-Freeness

Definition: A mechanism (X, t) is envy-free if

$c_i(X_i) + t_i \leq c_i(X_j) + t_j$ for any machines $i, j \in [m]$

Theorem [Cohen, Feldman, Fiat, Kaplan, Olonetsky, EC 2010]:
Any envy-free mechanism is at least
$\Omega(\log m / \log \log m)$-approximate.

Are there fairness notions that allow a
constant-factor approximation to optimal makespan?

Proportionality

Definition: A mechanism (X, t) is proportional if

$c_i(X_i) + t_i \leq (1/m) \cdot c_i([n])$ for any machine $i \in [m]$

	client 1	client 2	transfer	total cost	prop. share
salesman 1	2 hours	3 hours	½ hour	2½ hours	2½ hours
salesman 2	$1 - \epsilon$ hours	2 hours	−½ hour	1½ hours	$1½ - \epsilon/2$ hours

Salesman 2's total cost exceeds the proportionality threshold.

Proportionality

Main Theorem: There exists a proportional mechanism that achieves a $3/2$-approximation to makespan.

The $3/2$-approximation is tight.

Next slides:

- A 3/2 lower bound for any proportional mechanism.
- Characterization of proportionable allocations.
- A 3/2 upper bound via the Anti-Diagonal Mechanism.

Proportionality

Theorem: The 3/2-approximation ratio is tight.

Is the optimal allocation below **proportionable**?

	client 1	client 2	transfer	total cost	prop. share
 salesman 1	2 hours	3 hours	t hours	$2 + t$ hours	2½ hours
 salesman 2	$1 - \epsilon$ hours	2 hours	$-t$ hours	$2 - t$ hours	$1½ - \epsilon/2$ hours

No because: $(2 + t) + (2 - t) > 2½ + (1½ - \epsilon/2)$.

Characterization of Proportionality

Definition: An allocation is mean-efficient if

$$\sum_{i \in [m]} c_i(X_i) \leq \sum_{i \in [m]} \frac{c_i([n])}{m}$$

In words: The sum of processing times is at most the sum of proportional shares.

Theorem: An allocation is proportionable if and only if it is mean-efficient.

Anti-Diagonal Mechanism

Main Theorem: There exists a proportional mechanism that achieves a 3/2-approximation to makespan.

1. Assume WLOG the diagonal allocation minimizes makespan.
   (optimal makespan but may violate proportionability)
2. Choose the cost-minimizing anti-diagonal.
   (proportionable but may increase makespan)
3. Adjust allocation via merge and swap operations.
   (proportionable with near-optimal makespan)

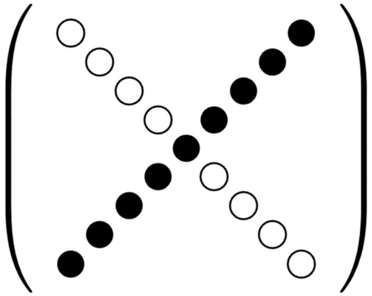

Anti-Diagonal Allocation

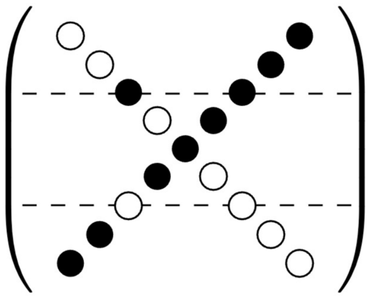

Merge Operation

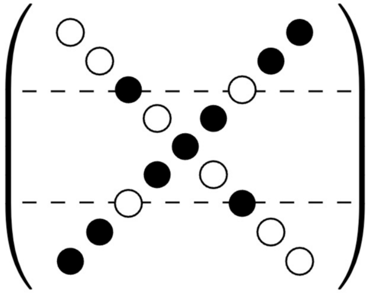

Swap Operation

Normalized Instances

Definition: An instance is **normalized** if $(1/m) \cdot c_i([n]) = 1$ for every machine i.

Theorem: There exists a **proportional** mechanism for normalized instances that achieves **optimal** makespan.

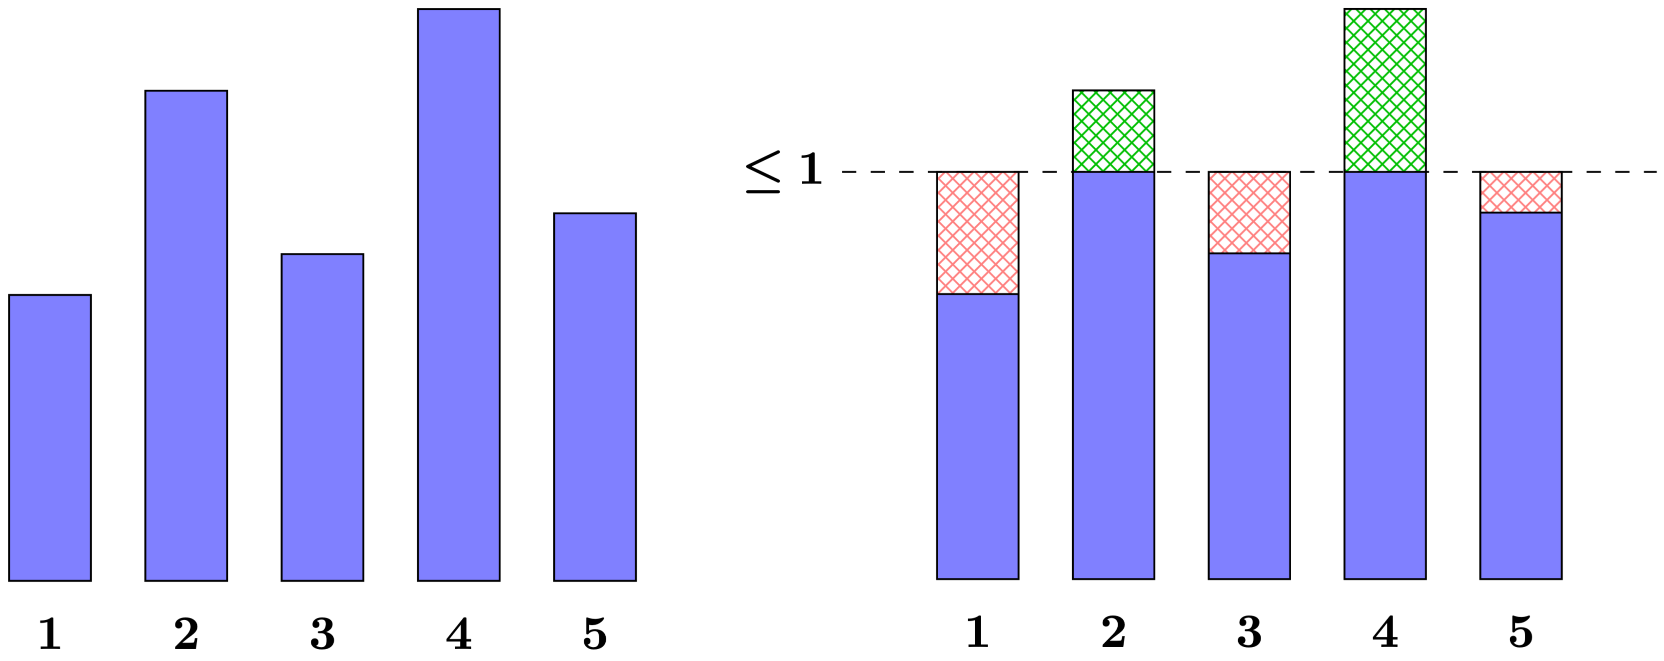

Conclusion

Proportionality allows a 3/2-approximation.
Envy-freeness requires a $\Omega(\log m \,/\, \log\log m)$-approx.

Can you achieve a constant-factor approximation of optimal makespan under alternative fairness notions?

Possibilities: EF1/EFX or $(1-\epsilon)$-EF with transfers.

The technique used to obtain the $\Omega(\log m \,/\, \log\log m)$ lower bound for EF does not extend to $(1-\epsilon)$-EF.

See the paper for details.