CSC2421: Online and Other Myopic Algorithms
Fall 2025

Allan Borodin

November 5, 2025

Week 9

Annoucements

- I am planning to start listing questions for the second assignment.
- Please submit your proposal for a project if you have not already done so.

Todays agenda

- Clarifying prirority algorithms vs polynomial time algorithms. Colouring d-Independent, chordal and d inductively independent graphs.
- The many dimensions of fair division.
- Review of the most studied fairness measures: EF (envy-freeness), Proportionality (prop) , MMS (max-min-share), MNW (max-Nash-welfare).
- Some results for EF, MMS, MNW
- If time permits, we return to Chapter 8 and extensions to the classical ski-rental, bin packing, and k-server problems.

Priority colouring and MIS for d-inductive and chordal graphs?

Here is where I misspoke in the October 22 class. I said the following:

"It is also the case that the PEO provides an optimal fixed order priority algorithm for computing a maximum independent set in a chordal graph (i.e., the MIS problem). For example (as we know), sorting by non decreasing finishing times, we obtain an optimal fixed order priority algoritihm for interval selection."

I was thinking in terms of the interval repesentation of an interval graph. BUT, it is not at all clear how one can can do this from a graph representation (e.g., the VAM-FI input model). One might still say we have a greedy algorithm for optimally coloring any chordal graph since we can create the PEO and the reverse of the PEO within polynomial time but this doesn't satisfy our definition of a priority algorithm

Priority algorithms for trees and bipartite graphs

When considering offline priority algorithms for graphs, the VAM-FI is the appropriate input model in contrast to online algorithms where VAM-PH is the more approriate model.

In the VAM-FI model, it is not difficult to see that we can optimally 2-colour bipartite graphs (and therefore trees) by an adaptive priority algorithm. We can simply start colouring any node r with color 1.

Then the next nodes in the ordering are the neighbours N_r of r in a breadth first search. The nodes in N_r receive color 2. The (not yet coloured) neighbours of each node v in N_r become the nodes next in the ordering, and they receive color 1. We continue in this way to colour any bipartite graph with 2 colours.

If we ever try to give conflicting colours to any node previously coloured than we know that the graph is not bipartite.

Creating d-inductive and d-inductively independent orderings

Can we create the orderings we need for classes of graphs that enjoy special orderings as fixed or adaptive priority orderings within the VAM-FI graph input model?

Note that every d-inductive graph is d-inductively independent and chordal graphs are 1-inductively independent.

The forward d-inductive and d-inductively independent orderings give us d-approximations for the MIS problem on these graphs.

The reverse d-inductive and reverse d-inductively independent orderings give us d-approximations for the colouring problem on these graphs.

It is easy to see the forward d-inductive ordering can be created as an adaptive priority ordering. We can simply determine the node with the smallest degree in the VAM-FI model.

Creating d-inductive and d-inductively independent orderings continued

But it doesn't seem possible to do this for a chordal graph (and hence for any d-inductively independent ordering in the VAM-FI model. That is, how can we determine how many neighbours of a vertex v are independent in the VAM-FI input model? However, we can create these forward orderings if we extend the VAM-FI model to also include all the nodes within distance 2.

But I do not see how we can create the reverse d-inductive ordering or the reverse PEO for chordal graphs (and hence the reverse ordering for d-inductively-independent graphs) as a priority orderings.

See Ye and Borodin [2012] for a discussion of d-inductively independent graphs, and the approximation bounds for the MIS problem and colouring problem for these graphs.

Fairness from the perspective of social choice theory and algorithms

The are many variants and issues concerning the “fair division” of goods and services. Indeed one can say that a central theme in political science is who should pay for and who should receive various services. We shall avoid the more controversial aspects of “fair decisions” (as in decisions for who gets loans, paroles, fair taxation, etc.) made by machine learning algorithms and focus on some precise well-studied meanings of “fairness”.

We will consider the division of both divisible and indivisible items but mostly focus on indivisible items. When we speak of *online* fair division, we can consider items or resources arriving sequentially, or agents arriving sequentially. In some problems, such as kidney transplant matching, it may be that both resources (e.g., donors) and agents (recipients) are arriving sequentially. In any (online or offline) fair division problem there are many dimensions to the problem. The following are some of the dimensions of fair division:

The main dimensions to fair-division

- What are the fairness criteria?
- Are decisions being made by a centralized mechanism or by decentralized self interested agents?
- Can randomized algorithms be considered fair?
- Are agents truthful?
- Are the items divisible or indivisible?
- Is there a single (divisible) item that is being shared or multiple (divisible or indivisible) items to be shared?
- And how do results change when we consider random order sequences or priority algorithms rather than online or arbitrary offline solutions?
- Is there any recourse to online decisions?
- Achieving better faincess approximations with predictions.

Fairness criteria

There are a number prominent fairness concepts including:

- Envy-freenes (EF): Agent i envies agent j if $v_i(A_j) > v_i(A_i)$. An allocation is envy-free if no agent envies another agent. One might say that "fairness is in the eye of the beholder". Some view EF as the "gold standard" for fairness. Unfortunately, as we will see, it is rarely obtainable. So we will have to settle for some weakening of EF.
- Proportionality: If agent i has value $v_i(S)$ when allocated the entire set S of items, then the allocation A_i that agnet i receives has value $v_i(A_i) \geq \frac{v_i(S)}{n}$.
- Max-Min fairness: The objective is to maximize the minimum allocation to any agent. The max-min objective has also been studied in combinatorial optimzation. In social choice theory it is often called *egalitarian social welfare*. In scheduling we consider the max-min allocation (rather than min-max scheduling in the makespan problem) to machines. Max-min scheduling has been called the Santa Claus problem. In scheduling it is not normally thought of as a fairness measure.

More fairness criteria

- Max-Min-Share (MMS). Agent i partitions S into $(A^i_1, A^i_2, \ldots, A^i_n)$. We define $MMS_i = \max_{A_1,\ldots,A_n}\{\min_j v_i(A_j)\}$. A partition $(A_1, \ldots, A_n)$ is an MMS partition (and hence considered fair) if $v_i(A_i) \geq MMS_i$ for all agents i. Why is this considered fair? Every agent reasons that there is a way to partition the goods for me to obtain my max-min share so I am entitled to at least this much value.
- The Nash Social Welfare NSW (or Nash Welfare) of an allocation $(A_1, A_2, \ldots, A_n)$ is defined as $(\Pi_{i=1}^n v_i(A_i))^{1/n}$. A Max Nash Welfare (MNW) solution is an allocation that maximizes the Nash Welfare over all possible allocations. Why is this a fairness condition?

More fairness criteria

- Max-Min-Share (MMS). Agent i partitions S into $(A^i_1, A^i_2, \ldots, A^i_n)$. We define $MMS_i = \max_{A_1,\ldots,A_n}\{\min_j v_i(A_j)\}$. A partition $(A_1, \ldots, A_n)$ is an MMS partition (and hence considered fair) if $v_i(A_i) \geq MMS_i$ for all agents i. Why is this considered fair? Every agent reasons that there is a way to partition the goods for me to obtain my max-min share so I am entitled to at least this much value.
- The Nash Social Welfare NSW (or Nash Welfare) of an allocation $(A_1, A_2, \ldots, A_n)$ is defined as $(\Pi_{i=1}^n v_i(A_i))^{1/n}$. A Max Nash Welfare (MNW) solution is an allocation that maximizes the Nash Welfare over all possible allocations. Why is this a fairness condition? A MNS solution implies some desireable fairness outcomes for indivisible goods including approximate MMS. Although, NP hard to compute the MNS, it can often be obtained in practice. See the publicly available site Spliddit (http //spliddit.org/) and the Caragiannis et al paper [2016,2019] where they state that the MNW solution is "arguably, the ultimate solution-for the division of indivisible goods". Note: Max-min $\leq$ MNW $\leq \sum_{i=1}^n v_i(A_i)/n$.

EF implies Proportional and Proportional implies MMS

- EF implies Proportional. Suppose an allocation $(A_1, \ldots, A_n)$ was not proportional. Then some agent i, $v_i(A_i) < v_i(S)/n$. This implies that some agent j has reeived more than $v_i(S)/n$ so that i envies j and hence the allocation was not EF.
- Proportional implies MMS. For any agent i, their least share is $\leq v_i(S)/n$

The previous inequalities apply to divisible or indivisible goods. We will mainly focus on indivisible objects and additive valuations. But we have to at least mention now what arguably might be the most popular fairness setting, namely cake cutting. As we probably all know, in cake cutting, we have n agents (e.g. birthday party attendees) that want their "fair share" of the cake. Cake cutting is clearly about a single divisible item and is usually stuied in the offline setting. with for example interesting results about how many cuts are needed to insure envy-freeness. Walsh [2011] defines and studies online cake cutting.

Indivisible goods, EF and EF1

We will assume n agents who are being allocated m indivisible items (goods). When we consider online allocations with the goods arriving online, it is more standard to say that we are allocating a sequence $t = 1, 2, \ldots T$ of online goods using T as the number of items. T may or may not be known to the mechanism or the agents.

It is easy to see that it is impossible to have an EF solution offline or online. Consider two agents who both want a single item. That item goes to one agent leaving the other agent envious.

As we previously indicated, we will have to settle for some weakening of EF to hope to obtain some degree of fairness.

Some weakenings of EF

The following provide some degree of envy freeness

- EF1 (*Envy freeness up to one good*): An allocation is EF1 if for all agents i and j, the exists an item $x \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{x\})$. That is, taking away some item from j's alloxcation removes the envy.
- EFX (*Envy freeness up to any good*): An allocation is EFX if for all agents i and j, and for all $x \in A_j$, $v_i(A_i) \geq v_i(A_j \setminus \{x\})$.
- Approximating EF: We can relax EF (and other fairness criteria) by allowing an approximation. That is, we may be quite satisfied if for all agents i and j, $v_i(A_i) \geq c \cdot v_i(A_j)$ for some approximation factor c that need not be a constant. For example, c might be a function of n or m. In many applications we expect that either n or m might be small so that such approximations can be useful.
- Bounded EF: There exists a constant b such that for all i and j, $v_i(A_i) \geq v_i(A_j) - b$.

Comments about EF1 and MMS

- EF1 is obtainable offline for all monotone valuations (i.e. $v_i(A_i) \leq v_i(A_i \cup \{x\})$; i.e., "free disposal". Lipton et al [2004] For additive valuations, the *round robin mechanism* obtains EF1.
- MMS cannot be obtained offline. Procaccia and Wang [2014] show that even for additive valuations, for $n \geq 3$ agents there cannot be a perfect MMS solution. Feige et al establish the current best $\frac{39}{40}$ inapproximation which is given by an example for 3 agents. The best offline approximation is $\frac{10}{13} \approx .7692$ by Heidari et al [2025] following the previous approximation of $\frac{3}{4} + \frac{3}{3836}$ by Akrami and Garg [2024].
- The situation gets much worse for online allocations even for additive valuations. With regard to online allocation of indivisible goods, when items arrive online both envy and MMS fairness have been studied. Online envy-freeness has been studied for randomized allocations. For MMS, I am only aware of deterministic algorithms. For allocating chores the situation is better in that there is a deterministic allocation of chores for any number of agents. There are also some results for MMS fairness when agents are arriving online.

A special case where EF1 can be obtained online

The following seems to be a folklore result that follows from the Lipton et al offline EF1 proof. .

Theorem: If all agents have $\{0, 1\}$ valuations, then there is an online algorithm that achieves EF1. (The offline EF1 result is based on cycle elimintaion in the envy-graph.)

Proof: As items arrive, we have an *envy graph* G_t at each time step where the nodes are agents and there is a directed edge (j, k) from j to k if agent j envies agent k.
For each time $t \geq 1$, we maintain the following invariant after the t^{th} item is allocated:

The allocation thus far is EF1 and the envy graph G_t is cycle free.

We prove by induction that the invariant holds for all t. The invariant is obviously true for $t = 1$. Assuming the invariant is true for $t - 1$, we show how to allocate the t^{th} item so that the invariant remains true.

Finishing the EF1 online algorithm for $\{0, 1\}$ valuations

We repeat the invariant and then provide the inductive step.

The allocation thus far is EF1 and the envy graph G_t is cycle free.

Since G is cycle free there must be a directed path from some node j to a node k where the in-degree of j is 0 and the out-degree of k is 0. Thus j is not envied by anyone. We allocate item t to agent j.

Since no one was envying j, giving the t^{th} item to j means that any other agent can envy j by at most one item. And since we have $\{0, 1\}$ valuations, and since the allocation was EF1 before the t^{th} item arrived, agent j no longer envies anyone and all the edges outgoing from j can be removed. Thus no cycle can be added.

It has been shown that for $\{0, 1\}$ valuations that MMS and EF1 are equivalent so that we obtain online perfect MMS for $\{0, 1\}$ valuations.

A randomized algorithm for EF

We consider what is arguably the simplest randomized allocation algorithm. The Like mechanism is introduced in Alexandrov et al [2015] and further studied in Alexandrov and Walsh [2019].

The Like algorithm randomly allocates an online item to any agent who states a positive value for the item; that is, if n_i agents have a posiitve value for the i^{th} item, and the j^{th} agent has a positive bid for item i, then the mechanism allocates item i to agent j with probability $\frac{1}{n_i}$.

We have some good and bad news reegarding the Like algorithm. First the good news.

- The Like algorithm is envy-free *ex-ante* for all **additive** valuations. That is, for every online sequence of items $g_1, g_2, \ldots, g_T$: $\mathbb{E}[Envy_T] = 0$. The result is not restricted to binary valuations. That is, for all agents j, k, $\mathbb{E}[v_j(A_j)] \geq \mathbb{E}[v_j(A_k)]$ and hence $\mathbb{E}[v_j(A_j)] = \mathbb{E}[v_j(A_k)]$.

More good news regading the Like mechanism

- The Like algorithm is strategy proof (SP) or truthful for all online sequences.
  Note: For online algorithms, one can consider a somewhat weaker form of strategy proofness; namely, an algorithm is online strategy proof (OSP) if for all items e_i, no agent j can strictly increase their expected utility by not reporting truthfully their valuation for e_i assuming full knowledge of the past decisions of all egents.
- In fact, any non-wasteful mechanism (i.e. greedy in the sense that it will always allocate an item to someone who has value for the item) that is strategy proof and envy-free ex ante is equivalent to the Like mechanism.
- Like achieves a $\frac{1}{n}$ approximation to Pareto efficiency ex-ante.

Other mechanisms for EF fairness

Alexandrov and Walsh [2019] introduce a number of online mechanisms and characterizations. They define variants of Like mechanisms including the following:

- Balanced Like: For an item i, amongst those agents j having $v_j(i) > 0$, allocate randomly to those who have received the fewest allocations thus far.
- Maximum Like: For an item i, amongst those agents j having $v_j(i) > 0$, allocate to an agent having the maximum value for that item.
- Online Random Priority: Choose a random ordering π of the agents initially. Then allocate item i to the highest ranked (wrt π) agent with $v_j(i) > 0$.

Balanced Like is not strategy proof; that is, after allocating all items, an agent might have wanted to be non-truthful at some earlier time. The next slide is a table from the Aleksandrov and Walsh paper summarizing properties of the 6 mechanisms discussed in their paper.

Properties of Fairness Allocaction algorithms in Aleksandrov and Walsh

Table 1. Axiomatic results. Key: $\star$ - the result follows from [Aleksandrov *et al.* 2015].

Mechanism	SP	OSP	EFA	SEFA	EFP	SEFP	BEFP	PEA	PEP
	General cardinal utilities								
Online RP	✓	✓	✓	✓	×	×	×	×	✓
Online SD	✓	✓	×	×	×	×	×	✓	✓
Maximum Like	×	×	×	×	×	×	×	✓	✓
Pareto Like	×	×	×	×	×	×	×	×	✓
Like	✓⋆	✓	✓⋆	✓	×⋆	×	×⋆	×	×
Balanced Like	×⋆	✓	×⋆	×	×⋆	×	×⋆	×	×
	Identical cardinal utilities								
Like	✓⋆	✓	✓⋆	✓	×⋆	×	×	✓	✓
Balanced Like	×	✓	✓	✓	×⋆	×	×	✓	✓
	Binary cardinal utilities								
Like	✓⋆	✓	✓⋆	✓	×⋆	×	×⋆	✓	✓
Balanced Like	×⋆	✓	✓⋆	×	×⋆	×	✓⋆	✓	✓

Note that Online Random Priority is Pareto Optimal ex-post but not Pareto Optimal ex-ante.

Almost matching good and bad news for randomized EF fairness allocations

We stated the Aleksandrov and Walsh definition for EF ex ante for which they show they show that the LIke mechanism is EF ex-ante. There is, however, a more meaningful quantifiable sense in which we can study randomized algorithms for envy freeness. Benade et al define $Envy_{j.k} = \max\{v_j(A_k) - v_j(A_j), 0\}$ and then study $\mathbb{E}[\max_{j,k}\{Envy_{j,k}\}]$. It can be seen that Aleksandrov and Walsh are studying $\max\{\mathbb{E}[Envy_{j,k}]\}$ which is 0 for the Like mechanism.

We have positive and negative results from Benade et al [2023,2025] where they consider envy regret over time. Consider a sequence of online items $e_1, e_2, \ldots, e_T$ and let $Envy_T = \mathbb{E}[\max_{j,k}\{Envy_{j,k}\}]$ at time T. They consider that an algorithm have vansihing regret if $Envy_T = o(T)$ so that $\frac{Envy_T}{T} \rightarrow 0$ as $T \rightarrow \infty$.

Results from Benade et al [2023,2025]

- For T sufficiently large ($T \geq n \log T$), the random allocation algorithm[1] achieves $\mathbb{E}[Envy_T] = O(\sqrt{TlogT/n}$ with respect an *adaptive* adversary. This implies the same bound for deterministic algorithms.
- For any $r < 1$, and $n \geq 2$, $\mathbb{E}[Envy_T = \Omega((\frac{T}{n})^{r/2})$. Hence we can say roughly that $\mathbb{E}[Envy_T] = \tilde{O}(\sqrt{T})$.
- No randomized algorithm can achieve both $\mathbb{E}[Envy_T] = o(T)$ and a $(\frac{1}{n} + \epsilon)$ approximation to Pareto efficiency ex-ante against an *oblvious* adversary.

[1]Benade et al seems to say that this algorithm randomly allocates to all agents uniformaly at random rather than the Like mechanism that only allocates to agents with positive value for an item. This sdeems to be irrelevant but it is curious that Benade et al do not refer explictly to the LIke mechanism.

A better bound for EF regret against an oblvious adversary

Halpern et al [2025] achieve a significantly improved regret bounds foir envy-freeneww (in the sense of Benade et al) for randomzed algorithms against an obvlivious adversary.

- There is a randomized algorithm with $Envy_T = O(\log T)$
- For every online randomized algorithm and every $r < \frac{1}{2}$, $Envy_T = \Omega(\log T)^r)$.
  Hence bound for $Envy_T$ is essentially optimal.

MMS fairness

When there are only 2 agents the classic cut and choose procedure achieves MMS fairness for both divisiible and indivisible items. However, as we have previously stated, when there are $n \geq 3$ agents, we can no longer achieve a perfect MMS even for additive valuations. The current best approximation in the offline setting achieves a $\frac{10}{13}$.

Perhaps, not surprisingly, the results for online MMS allocation are not as positive. We will continue the notation that there are n agents and the entire set of goods is S. Online MMS fairness has been considered for both online arrival of goods as well as online arrival of agents. In our discussion of MMS fairness we will only consider deterministic algorithms.

We begin by considering the online arrival of goods. We will assume that valuations are normalized so that $v_i(S) = n$. We note that it is sufficient to use the weaker assumption that $v_i(S)$ is known for all agents. We can observe when defining MMS fairness that $MMS_i \leq \frac{1}{n} v_i(S)$ and hence $MMS_i \leq 1$ assuming normalized valuations. In order to guarantee that MMS_i is an αMMS, it suffices to obtain an allocation for which $v_i(A_i) \geq \alpha$ for all agents i.

Online MMS fairness for 2 agents

When there are $n = 2$ agents, there is a deterministic online algorithm that allocates goods and achieves an $\frac{1}{2}$MMS allocation. Furthermore, $\frac{1}{2}$ is the best approximation to MMS that can be achieved for online goods when there are 2 agents. These results are from Zhou et al [2023]

We describe the algorithm and why it acheives the desired MMS approximation. Since it suffices for each agent to obtain a total value of $\frac{1}{2}$, once one of the two agents achieves that much value, the remaining goods should all go to the other agent. Until one of the agents achieves value at least $\frac{1}{2}$, we proceed as follows: If a new online item g has value at least $\frac{1}{2}$ for both agents, then allocate that item to the agent whose current value is the smallest amongst the two agents (breaking ties by allocating arbitrarily to agent 1). Otherwise allocate item g to the agent whose value for g is largest. (We break ties by allocating arbitrarily to agent 1.)

MMS fairness for $n = 2$ agents continued and bad news for $n \geq 3$

Note: I have to fill in more details for these resuts.

We cn prove that this algorithms achieves $\frac{1}{2}$ MMS fairness, proceeding by cases:
Case 1: Suppose an item g has value at least $\frac{1}{2}$ for both agents. **To be completed.**
Case 2: Suppose no item g has value at least $\frac{1}{2}$ for both agents. **To be completed.**

Theorem: For $n = 2$ agents, the best determinstic online approximation to MMS fairness is $\frac{1}{2}$. More precisely, for all $\delta > 0$, there is an input instance for which no deterministic online allocation can achieve a $\frac{1}{2} + \delta$ approximation to MMS fairness.
The following result is also from Zhou et al [2023].
Theorem
For $n \geq 3$, no deterministic online algorithm can achieve a constant approximation to MMS fairness.

MMS for Chores

We first give the definiton of MMS for chores. For chores, we have a cost $c_i(e)$ if agent i is allocated chore e. Costs are additive. Each agent wants to minmize the cost (i.e. work) for all the chores assigned to them.
As for goods, each agent chooses a partition.
$MMS_i = \min_{A_1,\ldots,A_n}\{\max_j c_i(A_j)\}$.
We will assume (like for goods) that costs are normalized so that $c_i(S) = n$ where S is the full set of chores. (As for goods, I think it is sufficient to assume that $v_i(S)$ is known for all agents.)

For chores $MMS_i \geq \frac{1}{n}c_i(S)$ and hence with nromaalization $MMS_i \geq 1$. Hence to achieve an $\alpha \geq 1$ MMS, it suffices to have an allocation such that every agent receives a total cost $c_i(S) \leq \alpha$. We also have the fact that $MMS_i \geq \max_{e \in S}\{c_i(e)\}$ since some bundle of goods must contain the chore of maximum value.

Allocating online chores continued

Theorem

For all $n \geq 2$, there is a deterministic algorithm such that the cost to each agent is at most $(2 - \frac{1}{n})$.

The algorithm is similar in spirit to the algorithm for goods with $n = 2$ agents. As soon as an agent has accumulated cost $\geq (1 - \frac{1}{n})$, that agent will not receive any more chores. For the agents with costs less than $(1 - \frac{1}{n})$, we assign each online chore "greedily" to the agent having minimum cost thus far.

For $n = 2$ agents, the previous algorithm gives a $\frac{3}{2}$ approximate MMS. There is a modified algorithm that achieves a $\sqrt{2}$ approximate MMS.

The current best negative results for chores are as follows (alsom from Zhou et al):

- For $n \geq 3$, the lower bound is 1.585 which is attained for the restriuctive case when all agents have identical valuations.
- For $n = 2$, the lower bound is $\frac{15}{11} \approx 1.634$

MMS fairness when agents are arriving online

We now turn to MMS fairness in the setting where each agent is arriving online to reflect applications such as disaster relief where agents (i.e. residents) request emergency goods. Here we assume that the algorithm only finds out about an agents valuation when the agent arrives online. Unfortunately, in this setting, the following example shows that no online algorithm can achieve a constant approximation to MMS fairness even for two agents with $\{0, 1\}$ valuations.

Suppose there are an even number m of items. Let the first arriving agent have value $= 1$ for each of the m items. This results in an MMS optimal partition with value $\frac{m}{2}$ by taking any subset of size $\frac{m}{2}$. To insure an allocation of $\alpha\frac{m}{2}$, the online algorithms awards her a bundle of size $\alpha\frac{m}{2}$ which are now no longer available to the second agent. The second agent has a value of 1 for exactly two items, say j_1 and j_2. If items j_1 and j_2 happen to be in the bundle given to agent one, every partition of the goods will result in a zero MMS value if $\alpha \geq \frac{2}{m}$.

Agents belonging to a small number of known types

When agents are arriving online, we have the following results for MMS fairness due to Kulkarni et al [2025].

We circumvent the negative consequence in the previous example by making information about the agents initially available to the algorithm. Specifically, we will assume that each agent i has a valuation $v_i : S \to \mathbb{R}^{\geq 0}$ belonging to one of k valuation types $\{v_j^*\}$ $(1 \leq j \leq k)$ initially known to the algorithm. If $v_i = v_\ell^*$, we say that agent i has type ℓ. When each agent arrives, they reveal their type. It is reasonable to assume that in many applications, such valuations are known and k is small.

Theorem:
There is a deterministic online algorithm ALG such that given an instance of k agent valuation types $v_j : S \to \mathbb{R}^{\geq 0}$, ALG achieves a $\frac{1}{k}$ approximation for MMS fairness.

MMS fairness with online agents continued

We complement the posiitve result for k types by the following lower bound.

Theorem:
For any deterministic online algorithm ALG_k, there is an instance with k valuation types such that ALG_k does not achieve an α MMS approximation with $\alpha \geq \frac{2}{\sqrt{k}-2}$.

We conclude the discussion of MM¡S fairness by asking which (if any) results can be improved with randomized algorithms (with respect to oblivious adversaries)? In particular, which negative results still hold?

Some concluding remarks on fairness

As we know, worst case adversarial results can be pessimistic, while stochastic assumptions (e.g., online inputs drawn i.i.d. for a known or unkownn distribution or random order online sequences) will often lead to much more positive results. In particular, lets consider envy-freeness in the i.i.d. model. That is, there is a known distribution D such that each online item is drawn independently from D. We end our discussion (for now) of fairness with the following result in Benade et al. :

For a known i.i.d. distribution, there is a randomized algorithm A computing an allocation such that for $T = T(\epsilon)$ (the number of online items in the input sequence) sufficiently large and all agents j, k, (1) Either j does not envy k up to one item or (2) Agent j does not envy agent k with probability at least $1 - \epsilon$.
In the latter case, we do not have a guarantee on how much agent j envies agent k (with probability at most ϵ). Hence we do not know if the regret bound for envy can be improved for i.i.d. online inputs.