CS 598RM: Algorithmic Game Theory, Fall 2020
HW 4 Solutions

---

1. (15 points)

   (a) (2.5 points) Consider a first price single item auction with set of bidders A. Let $n = |A|$, and private value v_i of each bidder i comes from a uniform distribution $[0,\ 1]$. Show that bid profile $(b_1, \ldots, b_n)$ such that $b_i = \frac{(n-1)}{n} v_i,\ \forall i \in A$ is a Nash equilibrium.
   Extra: What is the expected revenue of the auctioneer at this Nash equilibrium? How does it compare with the expected revenue of Vickrey Auction?

   (b) (2.5 points) Consider an arbitrary single-parameter environment with feasible set X. Prove that the welfare-maximizing allocation rule

$$x(b) = \text{argmax}_{x \in X} \sum_{i=1}^{n} b_i x_i \tag{1}$$

   is monotone.
   [For the definition of monotonic allocation rule, See Definition 4.2 of `http://theory.stanford.edu/~tim/f13/l/l3.pdf`]

   (c) (5 points) Consider a single item second-price (Vickrey) auction with n bidders. Assume that a subset S of the bidders decide to collude, that is, the bidders in S coordinate to submit false bids in order to maximize the sum of their payoffs. Under what conditions will this collusion be successful, and what should their collective strategy be?

   (d) (5 points) Consider a combinatorial auction in a which a bidder can submit multiple bids under different names, unbeknownst to the mechanism. The allocation and payment of a bidder is the union and sum of the allocations and payments, respectively assigned to all of its pseudonyms. Show that the following is possible: a bidder in a combinatorial auction can earn higher utility from the VCG mechanism by submitting multiple bids than by bidding truthfully.
   Can this ever happen in the Vickrey auction? Give a brief explanation.

   (a) Consider an agent i. We show that if every agent except i chooses the given bid, then i's BR to the strategy vector is $b_i = \frac{(n-1)}{n} v_i$. i's payoff for any bid b is $u_i = (b_i - v_i)Pr[b_i > b_j,\ \forall j \neq i]$. When the other agents bid the given value, this is $u_i = (b_i - v_i)Pr[b_i > \frac{(n-1)}{n} v_j,\ \forall j \neq i] = (b_i - v_i)\left(\frac{nb_i}{(n-1)}\right)^{(n-1)}$, where the last inequality follows as the values are uniformly distributed.
   We take the derivative of u_i with respect to b_i and equate it to zero to find the value of b_i that maximizes u_i.

$$u_i' = \left(\frac{nb_i}{(n-1)}\right)^{(n-1)} + n(b_i - v_i)\left(\frac{nb_i}{(n-1)}\right)^{(n-2)}.$$

$$u_i' = 0 \Rightarrow b_i = 0 \text{ OR } \frac{b_i}{(n-1)} = v_i - b_i \Rightarrow b_i = \frac{(n-1)}{n} v_i.$$

$b_i = 0$ minimizes the utility, and thus, the given bid, $\frac{(n-1)}{n} v_i$ maximizes i's utility hence is her BR.

(b) Fix an agent i. Let her allocation for bid vector $\mathbf{b}$ be x. Suppose i changes her bid to $z > b_i$, and the welfare maximizing allocation for the bid vector $(z, \mathbf{b_{-i}})$, is $\mathbf{y}$. $\mathbf{x}$ and $\mathbf{y}$ are welfare maximizing allocations for their corresponding bid vectors. Hence,

$$\sum_{j \neq i} b_j x_j + b_i x_i \geq \sum_{j \neq i} b_j y_j + b_i y_i, \text{ and}$$

$$\sum_{j \neq i} b_j x_j + z \cdot x_i \leq \sum_{j \neq i} b_j y_j + z \cdot y_i$$

Subtracting the first inequality from the second gives,

$$(z - b_i) x_i \leq (z - b_i) y_i \Rightarrow x_i \leq y_i,$$

as $z - b_i > 0$. Hence, if the bid of an agent increases, their allocation does not decrease. Thus the allocation rule is monotone.

(c) Let b_1, b_2 and b_3 be the bids of the first three highest bidders. Suppose the first and second highest bidders decide to collude. Then they can change their bids to some $\epsilon > 0$ greater than the third highest bid. The payoff to the highest bidder without collusion is $b_1 - b_2$, and that to the second highest bidder is 0. After colluding, the total payoff to the highest bidder is $b_1 - (b_3 + \epsilon) \geq b_1 - b_2$. The two bidders can decide to share the extra payoff $b_2 - (b_3 + \epsilon)$ in some way, thus getting at least as much payoff as received before colluding.

(d) We show an instance where the given scenario is possible.

Consider two agents with valuation functions denoted by u and v, and two items, say a and b. $u(a) = 1$, $u(a, b) = 1$, and $u(S) = 0$, for any set of items $S \neq \{a\}$. $v(a, b) = 1$, and $v(S) = 0$ for any set $S \neq \{a, b\}$.

The VCG auction can only satisfy one of the agents, and suppose gives both a, b to the agent with valuation u. The payments of the agents are $p_u = 1 - 0 = 1$, and $p_v = 1 - 1 = 0$, hence their utilities are 0.

Now suppose u enters a false bid with values $w(a, b) = 1$, and $w(S) = 0$ for every other set. The auction will allocate a to v and b to w, resulting in payments $p_u = 1 - 1 = 0$, $p_v = 2 - 2 = 0$ and $p_w = 1 - 1 = 0$. The utility of agent u is now $u(\{a, b\}) - p_1 - p_3 = 1$, strictly better than in the previous case.

This will not happen in a Vickery auction, as to change the outcome by submitting extra bids, one can only increase the second highest bid, thus decreasing the payoff of any agent.

2. (10 points) Combinatorial auctions. Recall the sponsored search auction problem discussed in class: there are k slots, the j^{th} slot has a known click-through rate (CTR) of α_j (non-increasing in j), and the payoff of bidder i in slot j is $\alpha_j(v_i - p_j)$, where v_i is the (private) value-per-click of the bidder and p_j is the price charged per-click in that slot. For historical reasons, modern search engines do not use the truthful auction discussed in class. Instead, they use auctions derived from the Generalized Second-Price (GSP) auction, defined as follows:

   (1) Rank advertisers by bid; assume without loss of generality that $b_1 \geq b_2 \geq \cdots \geq b_n$.

   (2) For $i = 1, 2, \ldots, k$, assign the i^{th} bidder to the i^{th} slot.

   (3) For $i = 1, 2, \ldots, k$, charge the i^{th} bidder a price of b_{i+1} per click.

   (a) (1 points) Prove that for $k = 2$ and any sequence $\alpha_1 \geq \alpha_2 > 0$, there exist valuations for the bidders such that the GSP auction is not truthful. (This in fact holds for any $k \geq 2$).

   (b) (2 points) Fix CTRs for slots and valuations-per-click for bidders. We can assume that $k = n$ by adding dummy slots with zero CTR (if $k < n$) or dummy bidders with zero valuation (if $k > n$). A bid vector b is an equilibrium of GSP if no bidder can increase its payoff by changing its bid. Verify that this translates to the following conditions, assuming that $b_1 \geq b_2 \geq \cdots \geq b_n$: for every i and higher slot $j < i$,

   $$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_j);$$

   and for every lower slot $j > i$,

   $$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_{j+1}).$$

   Hint: *Derive these by adopting i's perspective and "targeting" the slot j.*

   (c) (3 points) A bid vector $\mathbf{b}$ with $b_1 \geq \cdots \geq b_n$ is envy-free if for every bidder i and slot $j \neq i$,

   $$\alpha_i(v_i - p_i) \geq \alpha_j(v_i - p_j);$$

   where $p_j = b_{j+1}$

   Verify that an envy-free bid vector is necessarily an equilibrium. (The terminology "envy-free" stems from the following interpretation: for the current price-per-click is p_j of slot j; then the above inequalities say: "each bidder i is as happy getting its current slot at its current price as it would be getting any other slot and that slot's current price".)

   Hint: *You might want to first prove that the bidders must be sorted in non-increasing order of valuations.*

   (d) (4 points) Prove that, for every set of α_js and v_is, there is an equilibrium of the GSP auction for which the outcome (i.e., the assignment of bidders to slots) and the prices paid precisely match those of the truthful auction discussed in class. If you want, you can assume that the CTRs are strictly decreasing.

   Hint: *Recall that using Myerson's lemma we know a closed-form solution for the payments made by the truthful auction (see (8) and (9) in* `http://theory.stanford.edu/~tim/f13/l/l3.pdf`*). What bids would yield these payments in a GSP auction? Part (c) might be useful for proving that they form an equilibrium.*

(a) (James Hulett) The following valuation function leads to bidder 1′s best strategy being non-truthful, thus proving the claim. For slot 1, if $\alpha_1 > \alpha_2$, $\frac{\alpha_1}{(\alpha_1-\alpha_2)} > v_1 > 1$, and otherwise $v_1 > 1$, $v_2 = 1$ and $v_i = 0$ for all $i > 2$.

If agent 1 bids her true value, she will get assigned to slot 1, and pay the price α_1, getting utility $\alpha_1(v_1 - 1)$. If she lies and bids a value less than 1, she gets slot 2, pays 0, and gets utility $\alpha_2(v_1 - 0)$. As $v_1 < \frac{\alpha_1}{\alpha_1-\alpha_2}$ or $\alpha_1 = \alpha_2$, then $\alpha_2 v_1 > \alpha_1(v_1 - 1)$, hence her utility by lying is higher.

(b) Consider any agent i. To have that the current bid vector $[b_1, b_2, \ldots, b_n]$ is an equilibrium, we should have that b_i is i's BR to $\mathbf{b}_{-i}$. Now if i bids a value in (b_{i+1}, b_{i-1}), there will be no change in her payoff of $\alpha_i(v_i - b_{i+1})$, hence b_i is as good a bid as any of these. If i bids a value in (b_{j-1}, b_j) for some $j < i$, she gets assigned slot j and gets the payoff $\alpha_j(v_i - b_j)$. If i bids a value in (b_j, b_{j+1}) for some $j > i$, she gets assigned slot j and gets the payoff $\alpha_j(v_i - b_{j+1})$. Thus, for both scenarios to not give i a better payoff, we require,

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_j),\ \forall j < i,\ \text{and}$$

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_{j+1}), \forall j > i.$$

(c) We will show that the envy-free conditions imply the conditions from part b. Hence from the statement of part b, the bid vector is an equilibrium.

Fix an agent i. The left side of the envy-free condition is $\alpha_i(v_i - b_{i+1})$, the same as in the equilibrium conditions. The right side is $\alpha_j(v_i - b_{j+1})$. Hence for $j > i$, the envy-free condition is the same as the equilibrium condition. For $j < i$, as $b_j \geq b_{j+1}$, $\alpha_j(v_i - b_{j+1}) \geq \alpha_j(v_i - b_j)$. Hence, $\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_{j+1}) \geq \alpha_j(v_i - b_j)$, thus giving the equilibrium condition for all $j < i$.

Hence, an envy-free bid vector satisfies the equilibrium conditions from part b for any agent i, thus is an equilibrium.

(d) We will reverse-engineer the equilibrium from what conditions it is required to satisfy.

First, the prices paid must match those of the truthful auction. These prices have a well formed solution, given by,

$$p_i = \sum_{j=i}^{n} \frac{v_{j+1}}{\alpha_i}(\alpha_j - \alpha_{j+1}).$$

As the GSP auction has $p_j = b_{j+1}$ for every slot, we have that bid $b_j = p_{j-1}$. Thus, we get the bid vector by solving the equations for the prices iteratively. All that remains is to show this bid vector is an equilibrium. We will show the bid vector is envy-free, hence from part c, we will have it is an equilibrium.

We have $\alpha_i(v_i - p_i) = \alpha_i v_i - \sum\limits_{k=i}^{n} b_{k+1}(\alpha_k - \alpha_{k+1})$. Hence, we must show that for every $j \neq i$, $\alpha_i v_i - \sum\limits_{k=i}^{n} b_{k+1}(\alpha_k - \alpha_{k+1}) \geq \alpha_j v_i - \sum\limits_{k=j}^{n} b_{k+1}(\alpha_k - \alpha_{k+1})$. Call this equation (*).

For $j > i$, (*) simplifies to showing $(\alpha_i - \alpha_j)v_i \geq \sum_{k=i}^{j-1} b_{k+1}(\alpha_k - \alpha_{k+1})$. We prove this below.

$$\begin{aligned}(\alpha_i - \alpha_j)v_i &= v_i \sum_{k=i}^{j-1}(\alpha_k - \alpha_{k+1}) \\ &\geq \sum_{k=i}^{j-1} v_k(\alpha_k - \alpha_{k+1}) \text{ ...assuming } v_i \geq v_{i+1} \text{ for all } i \\ &\geq \sum_{k=i}^{j-1} b_k(\alpha_k - \alpha_{k+1}) \text{ ...as } v_i = b_i \geq b_{i+1} \text{ for all } i\end{aligned}$$

For $j < i$, (*) simplifies to $\sum_{k=j}^{i-1} b_{k+1}(\alpha_k - \alpha_{k+1}) \geq (\alpha_j - \alpha_i)v_i$. We prove this similarly as the previous case below.

$$\sum_{k=j}^{i-1} b_{k+1}(\alpha_k - \alpha_{k+1}) \geq b_i \sum_{k=j}^{i-1}(\alpha_k - \alpha_{k+1}) = v_i(\alpha_j - \alpha_i).$$

3. (5 points) Consider a combinatorial auction on m goods where you know a priori that every bidder is unit demand. This means that the valuation of a bidder i can be described by m private parameters (one per good) $v_{i1}, \ldots, v_{im}$, and its valuation for an arbitrary set S of goods is defined as $max_{j \in S} v_{ij}$. Prove that the VCG mechanism can be implemented in polynomial time for unit-demand bidders.

The allocation and payment rules of the VCG mechanism have been defined in the lecture notes. We will reduce the problem of computing the allocation and payments of the bidders to a polynomial number of max weight matching problems. As max weight matching is polynomial time solvable, the mechanism can be implemented in polynomial time.

Consider the complete bipartite graph G where bidders and items are the two parts, where every edge has weight v_{ij}, the value of bidder i for good j. Let G^{-i} be the same graph defined on all bidders except i.

The allocation rule of the VCG mechanism is to pick the welfare maximizing outcome. The welfare of an outcome is the sum of the highest valued good assigned to each bidder, i.e., $\mathbf{x}(\mathbf{b}) = \arg\max_{w \in \Omega} \sum_{i=1}^{n} \max_{j \in w_i} v_{ij}$.

Every max weight matching of G can be mapped to an allocation, by distributing the unmatched goods arbitrarily among the bidders and assigning the matched goods to the corresponding bidder. Note that no unmatched good is valued higher by any bidder than the one matched to them, as otherwise the weight of the matching could be improved, contradicting it being a max weight matching. Hence, the weight of the max weight matching is the welfare of the corresponding allocation.

Thus, finding a max weight matching is equivalent to finding a welfare maximizing allocation.

The payment rule of any bidder i can also be computed by finding the max weight matching of the graph G^{-i}, and taking the difference of the weight of this matching and the sum of weights of all bidders except i in G's max weight matching.

Hence, we can compute the allocations and payments of all bidders by finding $n + 1$ max weight matchings of $O(mn)$ sized graphs.

4. (10 points) Consider a set M of distinct items. There are n bidders, and each bidder i has a publicly known subset $T_i \subseteq M$ of items that it wants, and a private valuation v_i for getting them. IF bidder i is awarded a set S_i of items at a total price of p, then her utility is $v_i x_i - p$ where x_i is 1 if $T_i \subseteq S_i$ and 0 otherwise. This is a single-parameter environment. Since each item can only be awarded to one bidder, a subset W of bidders can all receive their desired subsets simultaneously if and only if $T_i \cap T_j = \emptyset$ for each distinct $i, j \in W$.

   (a) (3 points) Prove that the problem of computing a welfare-maximizing feasible outcome, given the v_i's and T_i's as input, is NP-hard.

   (b) (3 points) Here is a greedy algorithm for the social welfare maximization problem, given bids b from the bidders.

```
Initialize W = ∅ and X = M
Sort and re-index the bidders so that b_1 ≥ b_2 ≥ ··· ≥ b_n
for i=1,2,...,n do
    if T_i ⊆ X then
        remove T_i from X and add i to W
end for
Return winning bidders W.
```

   Does this algorithm define a monotone allocation rule? Prove it or give an explicit counterexample.

   (c) (4 points) If the above allocation rule is monotone then design a polynomial-time procedure/algorithm to compute payment of the winners. Otherwise do nothing.

(a) We reduce the known NP-Hard problem of finding a max weight independent set in a graph to the problem of computing a welfare maximizing feasible outcome of the given auction. To recall, the max weight independent set problem takes as input an undirected graph and asks to find a largest sized set of vertices such that no two vertices in the set are adjacent to each other (have an edge incident to both).

   Given the graph, we define the auction instance as follows. We have an agent for every vertex of the graph, and an item for every edge. The set T_i for every agent i is the set of edges adjacent to vertex i in G. The value v_i for T_i is 1.

   The welfare of a feasible outcome of the auction is the sum of agents who received their entire set T_i. For any allocation, we create an independent set in G by selecting the vertices whose agents received their T_i. This is an independent set as only one of any adjacent pair of agents can receive their common edge, hence no adjacent pair can receive their entire set T_i. Thus, a welfare maximizing allocation will correspond to the max size independent set.

(b) The algorithm does define a monotone allocation rule. To prove this, we show that for any bid vector $\mathbf{b}$, if an agent unilaterally deviates to a higher bid $b'_i > b_i$, she will get an allocation of value at least as much she got after bidding b_i.

   The algorithm orders bidders in non-increasing bid value. With b'_i, i's position in the order cannot be lower than her position with b_i. Hence, if she received her set T_i with bid b_i, no agent ranked 1 to $i-1$ was interested in any item in T_i. With a higher position, the agents now ranked before her are a subset of the previous agents, thus she will still get T_i. If she did not receive T_i with bid b_i, she got the lowest possible value, hence by bidding higher she cannot receive lesser.

Thus the allocation rule is monotone.

(c) As this is a single-parameter environment, we can use the payment rule specified in the proof of Myerson's lemma, and show it is computable in polynomial time.

Given a bid vector, there is exactly one *jump* in the value of any bidder i, at the point where she ranks higher than any other bidder i' with $T_{i'} \cap T_i \neq \emptyset$. Hence, the payment of a bidder is the bid value at which this jump occurs.

Given a bid vector, we compute the set of winners. For each winner i, we find the set of agents i' such that $T_{i'} \cap T_i \neq \emptyset$. We set $p_i = \max_{i'} b_{i'}$ as the payment of i. For all non-winning agents, we set their payment to 0.

Finding the set of winners and finding the set of agents i' for each winner i has a polynomial time brute force algorithm, hence the payment rule is computable in polynomial time.

5. (5 points) This problem considers a variation of the Bulow-Klemperer theorem (Nov 12 lecture). Consider selling $k \geq 1$ identical items (with at most one given to each bidder) to bidders with valuations drawn iid from F, where F is a regular distribution (i.e., corresponding ϕ^{-1} is monotonically increasing function). Prove that for every $n \geq k$, the expected revenue of the Vickrey auction (with no reserve) with $n + k$ bidders is at least that of the Myerson's optimal auction for F with n bidders.

(Hint: Myerson's optimal auction will be Vickrey with reserve $\phi^{-1}(0)$, *i.e.,* discard bids below $\phi^{-1}(0)$, give the item to k highest bidders and charge them $\max\{(k+1)^{th}\text{highest bid}, \phi^{-1}(0)\}$)

We define an auction, denoted by MV with $n + k$ bidders that runs in two steps as follows. MV runs Myerson's optimum auction on n bidders. All unsold items are distributed among k additional bidders at no cost. The second step does not generate any revenue, hence MV generates revenue equal to Myerson's auction.

We next show that the Vickery auction with $n + k$ bidders generates the highest revenue among all auctions where all items are sold. The feasible set is $X = \{x \mid \sum_i x_i = k\}$. A revenue maximizing truthful auction assumes $b_i = v_i$ and maximizes virtual welfar, i.e., $x^* \in \arg\max_{x \in X} \sum_i \phi(v_i)x_i$. Since F is regular, ϕ is monotone. Hence x^* assgins the items to k highest v_is. By Myerson's payment formula, the winners pay $b_{k+1} = v_{k+1}$. This is exactly the Vickery auction with k items.

Hence, the Vickery auction generates revenue equal or higher than any auction with $n + k$ bidders that sells all items, hence than MV too. As MV generates revenue equal to Myerson's auction, hence the Vickery auction generates revenue equal or higher than the Myerson's auction.

6. (5 points) Consider the reverse auction we briefly talked about in class: A denotes the set of bidders who are willing to sell the spectrum they hold. There is a set $F \subseteq 2^B$ of feasible sets that is upward closed (i.e., supersets of feasible sets are again feasible), i.e., if $S \in F$ then we can repack $A \setminus S$ in available range if spectrum of S is acquired.

   - Initialize $S = A$.
   - While there is a bidder $i \in S$ such that $S \setminus \{i\}$ is feasible:

     (*) Delete some such bidder i from S such that $S \setminus \{i\}$ is feasible.
   - Return S.

   Suppose we implement the step (*) using a scoring rule, which assigns a number to each bidder i. At each iteration, the bidder with the largest score (whose deletion does not destroy feasibility of S) gets deleted. The score assigned to a bidder i can depend on i's bid, the bids of other bidders that have already been deleted, the feasible set F, and the history of what happened in previous iterations. (Note a score is not allowed to depend on the value of the bids of other bidders that have not yet been deleted.) Assume that the scoring rule is increasing – holding everything fixed except for b_i, i's score is increasing in its bid b_i. Then, show that the allocation rule above is monotone: for every i and b_{-i}, if i wins with bid bid b_i then she will keep winning with any bid less than b_i.

Let $a_1, a_2 \ldots, a_n$ be the order in which agents in S got deleted with the bid vector $\mathbf{b}$. Consider an agent i. Suppose she unilaterally deviates to a lower bid $b_{i'}$. We must show that if i was the winning agent, she does not lose after bidding $b_{i'}$.

We show that the algorithm deletes agents in the same order $a_1 \ldots, a_n$, hence i wins again. For this, we show by induction on t that with the bid vector $(b_{i'}, \mathbf{b_{-i}})$, the agent with the highest score at any iteration t is a_t. For $t = 1$, the scores of all agents except i remain the same as with the bids $\mathbf{b}$, as they do not depend on the bids of agents who did not get deleted. i's score is increasing in its bid, hence can only become lower with $b_{i'}$. Hence, a_1 is the highest score agent again. Assuming this holds for the first k iterations, as the other bids do not change, the scores of the remaining agents in iteration $k+1$ are the same as with bids $\mathbf{b}$, as they do not depend on the bids of agents who did not get deleted and by inductive hypothesis, i is not deleted in the first k rounds. Hence, again as i's score can only become lower with $b_{i'}$, a_{k+1} is the highest score agent again.

Thus, agent i wins again.

7. (10 points) [Fair-division] Suppose there are n agents with additive valuations for m goods, such that EVERY AGENT'S MMS VALUE IS 1. If every good has value at most 1 for every agent, and there are at most n goods with value higher than 1/5 for any agent, then prove that there is a polynomial time algorithm to compute a 4/5-MMS allocation of the goods among the agents. (You can assume without loss of generality that the ordinal preferences of all agents are the same. See [BL16] for a proof.)

Let N be the set of all agents, and M the set of all goods.

Overview: We will first describe an efficient algorithm to get an EFX allocation A using the ordinal preferences are same assumption. Then if S is the set of agents who did not receive any good of value 1/5 or lower and G is the goods allocated to them by the EFX algorithm, we prove that the MMS value of every agent i in $(N \setminus G) \cup \{i\}$ for the goods $G \cup \{A_i\}$ is at least 1 (the same as its MMS value with N for M). Using the second observation, we sum the EFX property of every agent with agents in $N \setminus S$. This summed inequality gives that every agent's value for her own allocation is at least 4/5, hence the allocation is $4/5 - MMS$.

1. We prove that the following algorithm gives an EFX allocation.

Arrange the goods in non-decreasing order of ordinal preferences. Iteratively allocate goods to a source agent using the envy graph cycle elimination algorithm. Return the allocation obtained after allocating all goods.

We prove this allocation is EFX after t iterations for all t, by induction on t. For the base case, before the first iteration an empty allocation is EFX by default. Assume the allocation after allocating the first k goods is EFX. The $(k+1)^{st}$ good is allocated to a source agent. All new envy edges created are directed towards the source agent's allocation. No agent envied this allocation before allocating the last good. Hence, the envy is eliminated after removing the last good. The last good is the least valued good for every agent, as the ordinal preferences of all agents are the same and goods are being allocated in non-increasing preference order. Hence the allocation after the $(k+1)^{st}$ iteration is also EFX.

2. Next, note that if an agent does not receive a good of value 1/5 or lower, they receive at most one good. This is because there are at most n goods of higher value than 1/5, hence the EFX algorithm allocates these to distinct agents. Using this, we prove our next main observation.

3. The MMS value of every agent i for dividing goods in $G \cup A_i$ is at least 1.

Proof. Consider an MMS partition of M into $|N|$ bundles according to i's valuation function. It is given that the value of every bundle in this partition is at least 1. Now consider all bundles which contain goods in $M \setminus (G \cup A_i)$. Remove all goods from $G \cup A_i$ from these bundles, and distribute these arbitrarily among the remaining bundles. As there are $|N \setminus (S \cup \{i\})|$ items removed, there are at least $S \cup \{i\}$ bundles of value at least 1 remaining. This is an allocation of $G \cup A_i$ among $|S \cup \{i\}|$ agents with each bundle valued at least 1, hence i's MMS value is at least 1.

4. As the algorithm gives an EFX allocation, we have for every agent i,

$$v_i(A_i) \geq v_i(A_j) - v_i(g^{min}), \quad \forall j,$$

where g^{min} is the least valued item by i in A_j. For agents $j \in S$, $v_i(g^{min}) \leq 1/5$ for every i. Summing these inequalities over all agents in S gives,

$$\begin{aligned} v_i(A_i) &\geq \frac{1}{|S \cup \{i\}|} v_i(G \cup \{A_i\}) - \frac{1}{5} \\ &\geq MMM_i^{|S \cup \{i\}|}(G \cup \{A_i\}) - \frac{1}{5} \quad \text{...as MMS is at most average of all values} \\ &\geq 1 - \frac{1}{5} = \frac{4}{5} \quad \text{...as MMS is at least the original MMS.} \end{aligned}$$

Thus, every agent gets a bundle they value at least 4/5, hence the algorithm gives a 4/5-MMS allocation.

8. (10 points) Consider a game where n players are allocating a shared bandwidth of 1. Each player i chooses an amount $1 \geq x_i \geq 0$ and the utility of player i is $U_i(x_i, x_{-i}) = x_i(1 - \sum_{j=1}^{n} x_j)$. Note that if $\sum_{j=1}^{n} x_j > 1$, all players have a non-positive utility.

- (5 points) Does this game have a Nash equilibrium? If it does, give a Nash equilibrium.
- (5 points) Is it a potential game? Prove your answer.

- We prove that the strategy vector with $x_i = \frac{1}{n+1}$ for all i is a Nash equilibrium.
  The best response of a player i to the remaining strategy profile is obtained by taking the derivative of $U_i(x_i, x_{-i})$ with respect to x_i and equate it to zero to solve for x_i. $U_i'(x_i, x_{-i}) = (1 - \frac{n-1}{n+1} - x_i) - x_i = 0 \Rightarrow x_i = \frac{1}{n+1}$. Hence, no agent has incentive to deviate from their proposed strategy.
- The game is a potential game. We prove that $\phi(x) = \sum_i x_i - \sum_{i \leq j} x_i x_j$ is a valid potential function.

$$\begin{aligned}
&\phi(x_i', x_{-i}) - \phi(x_i, x_{-i}) \\
&= (x_i' + \sum_{j \neq i} x_j + \sum_{j \leq k, \neq i} x_j x_k + x_i' \sum_{j \neq i} x_j + x_i'^2) - (x_i + \sum_{j \neq i} x_j + \sum_{j \leq k, \neq i} x_j x_k + x_i \sum_{j} x_j) \\
&= x_i'(1 - \sum_{j \neq i} x_j - x_i') - x_i(1 - \sum_{j} x_j) \\
&= U_i(x_i', x_{-i}) - U_i(x_i, x_{-i}).
\end{aligned}$$

9. (10 points) Consider the following weighted generalization of the network cost-sharing game. For each player i, we have a weight $w_i > 0$. As before, each player selects a single path connecting her source and sink. But instead of sharing edge cost equally, players are now assigned cost shares in proportion to their weight. In particular, for a strategy vector S and edge e, let S_e denote those players whose path contains e, and let $W_e = \sum_{i \in S_e} w_i$ be the total weight of these players. Then player i pays $c_e w_i / W_e$ for each edge $e \in P_i$. Note that if all players have the same weight, this is the original game.

   - (4 points) Show that, in general, this game does not have an exact potential function.
   - (4 points) Show that there exists a potential function Φ such that,

$$\begin{array}{ll} \forall i, \forall \text{p}aths P_i, P_i': & (\Phi(P_i, P_{-i}) - \Phi(P_i', P_{-i}))(C_i(P_i, P_{-i}) - C_i(P_i', P_{-i})) \geq 0 \\ \forall i, \forall \text{p}aths P_i, P_i': & \Phi(P_i, P_{-i}) = \Phi(P_i', P_{-i}) \Leftrightarrow C_i(P_i, P_{-i}) = C_i(P_i', P_{-i}) \end{array}$$

   - (2 points) Using the above, show that the game has a pure Nash equilibrium.

- (Ben Pankow) Consider a network with 2 agents, say 1 and 2, one sink, one source, and 2 edges, say e_1 and e_2. Weights are $w_1 = 5$, $w_2 = 7$, $c_{e_1} = 1$, $c_{e_2} = 1$. Suppose for contradiction a potential function exists. Then we must have,

$$\phi(e_2, e_2) - \phi(e_2, e_1) = C_2(e_2, e_2) - C_2(e_2, e_1) = \frac{c_2 w_2}{w_1 + w_2} - \frac{c_1 w_2}{w_2} = \frac{7}{12} - 1$$
$$\phi(e_2, e_1) - \phi(e_1, e_1) = C_1(e_2, e_1) - C_1(e_1, e_1) = \frac{c_2 w_1}{w_1} - \frac{c_1 w_1}{w_1 + w_2} = 1 - \frac{5}{12}$$
$$\phi(e_2, e_2) - \phi(e_1, e_2) = C_1(e_2, e_2) - C_1(e_1, e_2) = \frac{c_2 w_1}{w_1 + w_2} - \frac{c_1 w_1}{w_1} = \frac{5}{12} - 1$$
$$\phi(e_1, e_2) - \phi(e_1, e_1) = C_2(e_1, e_2) - C_2(e_1, e_1) = \frac{c_2 w_2}{w_1} - \frac{c_1 w_2}{w_1 + w_2} = 1 - \frac{7}{12}$$

  We can express $\phi(e_2, e_2) - \phi(e_1, e_1)$ as both,

$$\phi(e_2, e_2) - \phi(e_1, e_1) = \phi(e_2, e_2) - \phi(e_2, e_1) + \phi(e_2, e_1) - \phi(e_1, e_1), \text{ and}$$
$$\phi(e_2, e_2) - \phi(e_1, e_1) = \phi(e_2, e_2) - \phi(e_1, e_2) + \phi(e_1, e_2) - \phi(e_1, e_1).$$

  Simplifying these using the previous equations we get,

$$\phi(e_2, e_2) - \phi(e_1, e_1) = \frac{7}{12} - 1 + 1 - \frac{5}{12} = \frac{2}{12}, \text{ and}$$
$$\phi(e_2, e_2) - \phi(e_1, e_1) = \frac{5}{12} - 1 + 1 - \frac{7}{12} = \frac{-2}{12}.$$

  Hence, such a potential function cannot exist.

- Extra credit
- Consider the strategy vector with the minimum value of $\phi(P)$. If it is not a PNE, then there is at least one agent who can unilaterally deviate to a path with lower cost. The value of $\phi(.)$ for the new strategy vector is also lower, from part b, a contradiction. Hence, this strategy vector is a pure Nash equilibrium.

References

[BL16] Sylvain Bouveret and Michel Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, 2016.