Online Ascending Auctions for Gradually Expiring Items (Extended Abstract)

Ron Lavi[*] Noam Nisan[†]

Abstract

In this paper we consider online auction mechanisms for the allocation of M items that are identical to each other except for the fact that they have different expiration times, and each item must be allocated before it expires. Players arrive at different times, and wish to buy one item before their deadline. The main difficulty is that players act "selfishly" and may mis-report their values, deadlines, or arrival times. We begin by showing that the usual notion of truthfulness (where players follow a single dominant strategy) cannot be used in this case, since any (deterministic) truthful auction cannot obtain better than an M-approximation of the social welfare. Therefore, instead of designing auctions in which players should follow a single strategy, we design two auctions that perform well under a wide class of selfish, "semi-myopic", strategies. For every combination of such strategies, the auction is associated with a different algorithm, and so we have a family of "semi-myopic" algorithms. We show that any algorithm in this family obtains a 3-approximation, and by this conclude that our auctions will perform well under any choice of such semi-myopic behaviors. We next turn to provide a game-theoretic justification for acting in such a semi-myopic way. We suggest a new notion of "Set-Nash" equilibrium, where we cannot pin-point a single best-response strategy, but rather only a set of possible best-response strategies. We show that our auctions have a Set-Nash equilibrium which is all semi-myopic, hence guarantees a 3-approximation. We believe that this notion is of independent interest.

*SISL/IST, California Institute of Technology. ronlavi@ist.caltech.edu. Supported by grants from the Israeli Science Foundation and the USA-Israel Bi-National Science Foundation. This work was done while this author was at the School of Computer Science and Engineering, The Hebrew University.

†School of Computer Science and Engineering, The Hebrew University. noam@cs.huji.ac.il. Supported by grants from the Israeli Science Foundation and the USA-Israel Bi-National Science Foundation.

1 Introduction

In recent years we have seen a growing body of work (e.g. [25, 1, 11, 12, 9, 23]) that analyzes distributed computer systems under the the assumption that participants will behave "selfishly" in the sense of optimizing their own utility, rather than behaving obediently or maliciously as was commonly considered in computer science. In this model, in order to design algorithms that will lead to desirable outcomes, one needs two ingredients. The first is a protocol, handed in to the participants. The second is a prediction of expected player behaviors (who tune their actions with respect to the given protocol). With these, we get a precise algorithmic description, and we are able to analyze its performance. In principle, for a given protocol, one may expect several different player behaviors. However, most recent works in this area manage to avoid this difficulty by designing truthful auctions. In a truthful auction, payments are devised in a sophisticated way so that players will maximize their utility by simply revealing their true input (i.e. follow a very specific, simple behavior). By this, the problem reduces to an algorithmic construction that should satisfy one more important requirement – truthfulness. The difficulty of analyzing several different player behaviors is therefore avoided.

Unfortunately, for many settings, truthful algorithms are rare, and a need to find other suitable solutions rises. In this paper we study a problem that forces us to take a different trail, instead of truthfulness. As we (provably) cannot design algorithms for which the players will be expected to take one single behavior, we design auctions for which *many* selfish behaviors lead to an approximately optimal allocation. Thus, our algorithmic construction is of a *family* of algorithms, each one corresponds to a specific combination of players' behaviors, and *all* of them obtain a near optimal outcome. The new concept is that, although players are not expected to follow a specific behavior, but only one out of a set of behaviors, the outcome is still guaranteed to be close to optimal, for any choice the players make. We also provide some game-theoretic rational why will the players limit their choice to this set of actions. We discuss some natural strengthenings of the equilibria notion we use, still keeping this general idea of "set equilibria". We believe that these concepts offer a new way to bypass the inherent difficulties of the truthfulness notion, in a way that suits the CS worst-case notions.

The problem we study is the online allocation of M items that are all identical except that they "expire" at different times: the first item expires at time 1, the second at time 2, and so on. Players arrive over time, and items must be allocated at or before their expiration time. Each player j desires any single item between his arrival time, r_j, and his deadline, d_j, and has a value v_j for receiving the item. All information r_j, d_j, v_j is private to player j, and players act rationally to maximize their utility: the value v_j, if they are allocated an item, minus any payment that they must pay. Our goal is to design a mechanism that maximizes the "social welfare", i.e. to allocate the items so that the sum of values of players that receive an item is maximized.

This model seems applicable to many scenarios in which items are sequentially allocated as time progresses, where both items and players have a finite "lifetime". In a computational setting, this model is equivalent to online scheduling of unit length jobs with deadlines. Focusing on the algorithmic question only, and ignoring incentive issues, it is known that the offline problem can be solved exactly in polynomial time and that the online problem has a simple greedy algorithm that achieves a 2-approximation [18, 4], but no online algorithm can achieve an approximation ratio better than the "golden ratio" [15]. Of course, this algorithm requires each player to reveal his true type (value, arrival time, and deadline) and ignores players' strategic considerations.

To incorporate the strategic considerations of the players, our first attempt was to design a truthful mechanism for this problem, in which players are motivated to reveal their true input by incorporating some sophisticated payment scheme. However, this cannot be achieved, as the following strong impossibility result shows:

Theorem: *Any truthful deterministic online mechanism cannot obtain an approximation ratio* $< M$.

One could approach this difficulty by adding more assumptions about the players. E.g., a common assumption in recent works on online auctions [20, 8, 3, 5, 7, 13, 2, 19, 26, 16] is that player values are taken from some known interval $[v_{min}, v_{max}]$. With this, one can construct a randomized truthful auction with an approximation ratio of $O(log(v_{max} - v_{min}))$ (this can be obtained as a special case of the market clearing algorithm of [8], or by using the general method of [2] to convert online algorithms to truthful online algorithms). To our view, this is a too heavy toll to pay for truthfulness, as a deterministic 2-approximation without any assumptions on players exists once truthfulness is dropped. In addition, the resulting truthful auctions sometimes appear somewhat artificial (e.g. without real competition among the different bidders – all items are sold in a fixed price, determined before the first bidder arrives).

Instead, we will not assume any additional assumptions on players, but will relax the required notion of equilibria: instead of specifying a single tuple of strategies (the equilibrium point) that provides a good approximation ratio, we will specify a large set of strategies with the property that the mechanism will perform well on *any* of them. Our strategic analysis will not be able to pinpoint exactly which of these strategies will be rationally chosen, but rather only that one of them will be – this is enough to guarantee good performance.

At this point we embark with the algorithmic analysis of two (variants of) classic ascending auctions, under a wide class of possible player behaviors. The first auction we consider is a natural adaptation of the iterative auction of Demange, Gale, and Sotomayor [10] (similar offline scheduling auctions were also considered by [27]). *The Online Iterative Auction* constantly maintains a current price p_t and a current winner win_t for every item t. At each time t, each player (in his turn) may place his name as the temporary winner of some item t' (bid on t'), deleting the previous temporary winner, and increasing the price by some fixed small δ (a player can be a temporary winner only for one item). When none of the players wishes to bid, the time t phase ends: item t is sold to player win_t for a price of $p_t - \delta$. At time $t+1$ the prices and temporary winners from time t are kept, and the auction continues similarly.

In the offline setting, where all players arrive at time 1, [10] show that if all players behave *myopically*, i.e. always bid on the item with the lowest price among those that interest them, then the auction will reach the optimal allocation. Moreover, such myopic behavior is indeed the player's best interest [14]. But what will the player choose, facing this auction in the online setting? This depends on the beliefs of the player about the future: if he fears that new competitive bidders will arrive in the future, he may bid aggressively for earlier items, offering a higher price for them but reducing his risk of future competition. To incorporate such considerations, we call a player *semi-myopic* if he always bids on some item with price lower than his value (not necessarily the item with the lowest price, as the myopic behavior requires). Thus there exist many semi-myopic behaviors, that represent different beliefs. The point is that the auction will obtain near optimal allocation under *any* combination of such behaviors:

Theorem: *If all players are semi-myopic then the Online Iterative Auction achieves a 3-approximation of the welfare.*

We prove this by analyzing a family of semi-myopic algorithms, where each such algorithm corresponds to a specific combination of semi-myopic behaviors. One algorithm in the family is the greedy algorithm[1], but the analysis of the entire family is completely different and non-trivial (even the analysis of the "myopic algorithm", which results from the myopic behaviors of the players, is completely different).

The second auction we consider is *The Sequential Japanese Auction*: Item t is sold at time t using a classic one-item ascending auction (exact details appear in the paper body). Surprisingly, we show that this auction has a similar structure to the previous one (in our setting). We define a myopic behavior that leads to the optimal allocation in the offline case (when all players arrive at time 1), and, similarly to above, a family of semi-myopic behaviors aimed to capture players' uncertainties about the future. These semi-myopic behaviors again exactly correspond to our family of semi-myopic algorithms, hence a 3-approximation is obtained for every combination of semi-myopic behaviors.

But why should the players play as we expect? We now turn to give a more accurate game-theoretic analysis of the players' behaviors. As truthful auctions do not exist, we instead seek an equilibrium notion that will capture the idea advocated above, i.e. that the best we can do is recommend on a set of strategies, and not on a specific, single strategy.

In the game-theoretic setting, each player is required to choose a strategy $s_i \in S_i$. The resulting payoff of each player i is $u_i(s_1 ... s_n)$[2]. Player i's strategy $r_i \in S_i$ is a best response to a specific combination of strategies of the other players $s_{-i} \in S_{-i}$ if for any $s_i \in S_i$, $u_i(r_i, s_{-i}) \geq u_i(s_i, s_{-i})$. Our notion of "set equilibria" captures the situation where we describe a subset $R_i \subseteq S_i$ of "recommended strategies" (best response strategies) to choose from, instead of describing a single strategy r_i as the equilibrium point:

[1]which corresponds to the completely aggressive behavior that continues to bid on the current item as long as its price is lower than the player's value (even if the next item has a price of zero).

[2]We actually need the framework of games with incomplete information, in which each player additionally has a type (in our case, a triplet {arrival time, value, and deadline}), and both his strategy and his payoff may depend on this type; this adds few technicalities that we deal with in the paper body, but does not change the spirit of the results, as we describe here.

Definition: *The sets R_i are in* **Set-Nash equilibrium** *if for any player i, and any strategy combination of the other players $s_{-i} \in R_{-i}$, player i has a best response to s_{-i} in R_i.*

Thus the definition requires that a best response to any tuple of recommended strategies of the others may be found within the recommended strategies of player i. This becomes equivalent to regular Nash equilibrium when $|R_i| = 1$ for all i. It should be pointed out that there always exists a trivial Set-Nash equilibrium in which the recommended strategies are the entire set of strategies. Therefore this notion is interesting only when one can guarantee some performance bound whenever players play any one of their recommended strategies, as we do.

Although Set-Nash equilibrium is a weak notion, certainly weaker than regular Nash equilibrium, it seems to us that it carry some weight, especially in computerized environments, in which appropriate protocols and software programs that act "as recommended" are available, and so a deviation would seem to require some effort. In such cases players can be realistically expected to act as recommended unless they have clear incentives to deviate. Such a clear incentive would seem to be absent when the recommended strategies are in Set-Nash equilibrium.

We also provide some discussion on ways to strengthen the basic definition. We describe a hierarchy of four "set equilibria" notions, with a growing strength. While, for our motivating example, we were able to use only the basic definition, we believe that the complete hierarchy will turn out to be useful for other models, in which truthfulness does not exist, and one wishes to remain within the worst-case framework and to avoid any strong distributional assumptions.

Returning to our model, we show that both our online ascending auctions have Set-Nash equilibria that are all semi-myopic. We leave the description of the appropriate sets of recommended strategies to the body of the paper. The main point we arrive at is that players do not have a clear incentive to deviate outside of these sets of recommended strategies; and when they do stay inside the set of recommended strategies, the mechanism obtains a 3-approximation.

Main Theorem: *The Online Iterative Auction and the Sequential Japanese Auction both have a Set-Nash equilibrium which is all semi-myopic, hence results in a 3-approximation of the welfare.*

The rest of the paper is organized as follows. The model and basic definitions are given in section 2. In section 3 we describe the two online ascending auctions, and show their algorithmic properties by characterizing a family of 3-approximations. Section 4 returns to the strategic setting, showing that no truthful auction can achieve an approximation ratio better than M. In section 5 we define the new notion of Set-Nash equilibrium, and give an exposition to the analysis of our auctions according to it. Full proofs and additional details appear in the full paper [21].

2 Model and Basic Definitions

Items: We wish to sell M identical items with different expiration times. W.l.o.g. we assume that the first item expires at time 1, the second at time 2, and so on. Each item must be sold (and received by the buyer) at or before its expiration time.

Players: The potential buyers of the items (players/bidders) arrive over time. Player i arrives to the market at time $r(i)$, and stays in the market for some fixed period of time, until his deadline $d(i)$. We assume w.l.o.g. that the arrival and departure times are integers[3]. Each player desires only one item (unit demand), that expires no earlier than his arrival time. He must receive it at or before his departure time[4]. Player i obtains a value of $v(i)$ from receiving such an item, otherwise his value is 0. We assume w.l.o.g. that different players have different values[5].

We assume the standard game-theoretic setting: Player i privately obtains his variables $r(i)$, $d(i)$, and $v(i)$, at time $r(i)$. He acts selfishly in order to maximize his own utility: his obtained value minus his price. I.e., a player may arrive at or after his true arrival time, and declare or act as if he has any value, and any deadline.

Our goal: is to maximize the social welfare: the sum of (true) values of players that receive an item.

Basic definitions: Player i is *active* at time t if $r(i) \leq t \leq d(i)$, and i did not win any item before time t. Let A_t be the set of all active players at time t. An *allocation* is a mapping of items to players such that, if player i receives item t, then $r(i) \leq t \leq d(i)$. Let X_t be an allocation of items $t, ..., M$. $X_t[d]$ denotes the player that receives item d according to X_t, and $X_t[d_1, d_2] = \cup_{d=d_1}^{d_2} X_t[d]$, the set of players that receive items d_1 through d_2. By a slight abuse of notation we also use X_t as the set of players $X_t[t, M]$. The *value* of X_t is $v(X_t) = \sum_{d=t}^{M} v(X_t[d])$, i.e. the welfare obtained by X_t. A set S of players is *independent* with respect to items $t, ..., M$ if there exists an allocation of (part of) the items $t, ..., M$ s.t. every player in S receives an item.

The offline allocation problem: The offline problem, in which all players arrive at time 1, is a matroid: a set of *players* is independent if there exists an allocation of (part of the items) to these players. This is known [17] for the unit-demand scheduling problem, which is equivalent to ours. This matroid structure is used extensively in our proofs. See the full paper for details [21].

3 Two Online Ascending Auctions

We first describe online adaptations of two well-known ascending auctions. These have the property that players do not have to choose specific actions for the auction to perform well: a 3-approximation is obtained for a large, reasonable family of behaviors that we term "semi-myopic". Under any such player behaviors, each of our auctions belongs to a general family semi-myopic algorithms, that we characterize. We then show that any semi-myopic algorithm obtains a 3-approximation, and therefore conclude that our auctions lead to a near optimal allocation for any choice of semi myopic behaviors of the players.

In this section, we focus on the algorithmic side. Therefore we give only intuitive justifications for the player behaviors that we assume. For the same reason, we also omit few technicalities about prices and tie-breaking rules from the definitions. These are detailed when we analyze the strategic properties of our auctions.

3.1 The Online Iterative Auction

We consider an online adaptation of the iterative auction of Demange, Gale, and Sotomayor [10]:

DEFINITION 3.1. (THE ONLINE ITERATIVE AUCTION) *The Online Iterative Auction constantly maintains a current price p_t and a current winner win_t for every item t. These are initialized to zero at $t = 0$, and updated according to players' actions at each time t, as follows:*

- *Each player, in his turn, may place his name as the temporary winner of some item t', causing the previous winner to be deleted, and the price to increase by some fixed small δ. A player cannot perform this action, and must relinquish his turn, if he is already a temporary winner.*
- *When none of the players that are not temporary winners wishes to place their names somewhere, the time t phase ends: item t is sold to the player win_t for a price of $p_t - \delta$.*

[3] as actions in a non-integral time point can be deferred to the next integral point with no affect.

[4] Our auctions also fit the more severe restriction that player i cannot get an item $t > d(i)$. E.g., player i cannot attend Saturday's show if he is leaving on Friday, even if he receives the ticket before Friday.

[5] I.e. fix some arbitrary order over players, and set $v(i) \succ v(j)$ iff $v(i) > v(j)$ or $v(i) = v(j)$ and $i \succ j$.

- *At time $t+1$ the prices and temporary winners from time t are kept. If additional players arrive then the auction continues according to the above rules.*

Before analyzing the online auction, it is useful to take a glimpse at the offline case, in which all players arrive at time 1. This is a special case of the unit-demand model studied by [10], [14]:

DEFINITION 3.2. ([10]) *Player i has a* **myopic strategy** *in the iterative auction if, in his turn, he always places his name on the item $t \leq d(i)$ with the minimal price, unless the minimal price $\geq v(i)$, in which case he does not place his name at all.*

LEMMA 3.1. ([10], [14]) *If all players are myopic and arrive at time 1 then the online iterative auction obtains the optimal allocation. Furthermore, if all other players are myopic then player i will maximize his utility by playing myopically.*

In the online setting, however, a player might not be completely myopic, depending on his beliefs about the future. For example, he may bid aggressively for the current item, not placing his name on future items at all. This is reasonable if he anticipates tight competition from players that will arrive later on. Viewing this behavior as one extreme, and the completely myopic behavior as the other, it seems that any combination of the two cannot be "ruled-out". On the other hand, a player might choose not to participate at all for some time units – if, for example, there are M high valued players that desire any item 1 through M, but they all do not participate up to time M, then the resulting welfare will be low. As it turns out, this is the only type of behavior we need to exclude:

DEFINITION 3.3. *Player i is* **semi-myopic** *if, in his turn, i places his name on some item t with $p(t) \leq v(i)$ and $r(i) \leq t \leq d(i)$ (not necessarily the one with the lowest price). If there is no such item, i stops participating.*

THEOREM 3.1. *If all players are semi-myopic then the online iterative auction achieves almost a 3-approximation: $v(OPT) \leq 3 \cdot v(ON) + 2 \cdot M \cdot \delta$, where OPT, ON are the optimal, online allocations.*

The proof shows that, under any semi-myopic behavior, the online iterative auction is a semi myopic algorithm (see section 3.3 below), hence obtains the desired approximation (by lemma 3.3).

3.2 The Sequential Japanese Auction

A different possibility is to sell item t at time t using a simple one item ascending auction:

DEFINITION 3.4. (A JAPANESE AUCTION) *The (classic, one item) Japanese auction operates as follows: An auctioneer gradually raises a price, starting from 0. Each participating player should decide whether to drop out or to stay (once a player drops out, he cannot join again), as the price ascends. The price stops increasing exactly when all players, besides one, have dropped out. The winner is the player that did not drop out, and he pays the price that was reached.*

A natural adaptation of this to the online case is:

DEFINITION 3.5. (THE SEQUENTIAL JAPANESE AUCTION) *The Sequential Japanese Auction sells each item t at time t, separately, using a Japanese auction with one modification: the participants are allowed to observe how many drop-outs occur as the price ascends (and to incorporate this into their drop-out decision).* [6]

As before, it is useful to first consider this auction for the offline case, in which a rather surprising notion of myopic behavior leads to the optimal allocation:

DEFINITION 3.6. *Player i is* **myopic** *in the Sequential Japanese Auction if, in the auction of any time t, (for $r(i) \leq t \leq d(i)$), he drops exactly when either the price reaches $v(i)$, or when there are exactly $d(i) - t$ other players that did not drop yet.*

The logic for dropping when $d(i) - t$ players remain is that at this point the player is assured that there are enough items before his deadline to be allocated to all bidders who are willing to pay the current price.

LEMMA 3.2. *If all players are myopic and arrive at time 1 then the Sequential Japanese Auction obtains the optimal allocation.*

Our assumption that player have different values is important here. It is not hard to verify that this lemma actually follows from the proof of the online strategic setting. In this case, a myopic behavior (in the offline case) is a best response when all others are myopic only when using the modified prices of this auction (described in the full paper).

In the online setting, again, players might not play myopically, and may insist on closer items (i.e. stay longer in the auction) if they anticipate much competition in the future. In the extreme, when every player remains in the auction until the price reaches his true value, we actually simulate the simple greedy

[6] Prices are also modified. The time-t-winner pays the highest price among all time-t'-auctions in which he tied the time-t'-winner. Defining "a tie" is delicate, and requires the players to drop simultaneously. See the full paper.

algorithm, which is a 2-approximation. As before, any behavior between the two extremes can cause only a minor performance degradation. All we wish is that players will not drop out "too soon". Indeed, dropping out early in the auction also have disadvantages, as future auctions might be much more competitive, due to new arriving players.

DEFINITION 3.7. *Player i's strategy is* **semi-myopic** *(for the Sequential Japanese Auction) if, at every time t, he drops no later than when the price reaches his value, $v(i)$, and no earlier than when only $d(i)-t$ other players remain in the auction.*

THEOREM 3.2. *If all players play semi-myopic strategies then the Sequential Japanese Auction achieves a 3-approximation.*

In a similar manner to the iterative auction above, this theorem is proved by showing that, under any semi-myopic behavior, the Sequential Japanese Auction is a semi myopic algorithm.

3.3 Semi-Myopic Algorithms For each combination of player strategies, the above auctions are associated with a different algorithm. In order to analyze their performance for a family of strategies, we therefore need to characterize a family of algorithms, that we call semi-myopic algorithms. The main point is that *any* semi myopic algorithm obtains a 3-approximation of the welfare.

Specifically, the *current best schedule* at time t, S_t, is the allocation with maximal value among all allocations of items $t, ..., M$ to the active players, A_t[7]. Define:

$$(3.1)\quad f_t = \{\ j \in S_t \mid S_t \setminus j \text{ is independent w.r.t items } t+1, ..., M\ \},$$

The set f_t contains all players that can receive item t, when one plans to allocate items $t, ..., M$ to the players of S_t (i.e. these are all the potentially *first* players). Now define the critical value at time t, v_t^*, as:

$$v_t^* = \begin{cases} \min_{j \in f_t}\{v(j)\} & S_t \neq f_t \\ 0 & \text{otherwise} \end{cases}$$

All active players with value larger than v_t^* must belong to S_t, because of its optimality (w.l.o.g the first player in S_t has value v_t^*, and if there was a higher valued player outside of S_t, we could switch between them and increase the value of S_t). Thus, it seems reasonable not to allocate item t to a player with value less than v_t^*, as this player cannot belong to any optimal allocation. Surprisingly, this condition is enough to obtain approximately optimal allocations:

DEFINITION 3.8. (A SEMI MYOPIC ALGORITHM) *An algorithm is semi myopic if every item t is sold at time t to some player j with $v(j) \geq v_t^*$.* [8]

LEMMA 3.3. *The Online Iterative Auction with semi-myopic players and the Sequential Japanese Auction with semi-myopic players are both semi myopic algorithms.*

The family of semi-myopic algorithms can be viewed as the entire range between the following two extremes: the first is the greedy algorithm that always chooses the player with maximal value[9], and the second is the "myopic" algorithm that always chooses the player that determined v_t^*. These two extremes are 2-approximations (both were studied in the context of online scheduling [18, 6]). The entire family has only a slightly larger approximation ratio:

THEOREM 3.3. *Any semi-myopic algorithm is a 3-approximation of the welfare, and this is tight.*

Proof. We will show that any online allocation algorithm that produces an allocation ON has $v(OPT \setminus ON) \leq 2\sum_{t=1}^{M} v_t^*$, where OPT is the optimal allocation. From this, the theorem will follow immediately, as any semi-myopic algorithm has $v(ON[t]) \geq v_t^*$, and therefore $v(OPT) = v(OPT \setminus ON) + V(ON) \leq 2\sum_{t=1}^{M} v_t^* + v(ON) \leq 2 \cdot v(ON) + v(ON) = 3 \cdot v(ON)$.

We first state two useful claims about the structure of offline allocations.

PROPOSITION 3.1. *Let A, B be sets of players, where $A \subset B$. Let S_A, S_B be the allocation with optimal value for A, B, respectively (both are over the same set of items). Then if $j \in A$ but $j \notin S_A$ then $j \notin S_B$.*

PROPOSITION 3.2. *Let S be the allocation with maximal value over the set of players A and the set of items $t, ..., M$. Assume that S is not independent w.r.t items $t+1, ..., M$. Let $j \in S$ be the player with minimal value such that $S \setminus j$ is independent w.r.t items $t+1, ..., M$. Then $S \setminus j$ has maximal value among all independent sets w.r.t items $t+1, ..., M$ and players in A.*

[7] There exists one such allocation, by the matroid structure, and since different players have different values.

[8] A worst-case approximation cannot sell item t before time t, as a player with high value only for t may appear.

[9] Interestingly, this is a special case of the greedy algorithm of [22] for combinatorial auctions with sub-modular valuations. They study the offline case, but it is easy to verify that their algorithm actually works online.

We can now prove the theorem. Fix some scenario, and let OPT and ON be the optimal and online allocations for this scenario. We describe $f : OPT \setminus ON \rightarrow \{1, ..., M\}$ such that f is 2 to 1 and $v(j) \leq v^*_{f(j)}$ for any $j \in OPT \setminus ON$. Our claim immediately follows. The function f is defined as follows. Let X_t be the optimal allocation of items $t+1, ..., M$ among players in $OPT[1,t] \setminus ON$. For any $j \in OPT \setminus ON$ (say $j = OPT[t']$), let $t^*_j = min\{t \geq t' \mid j \notin X_t\}$. Then we fix $f(j) = t^*_j$.

PROPOSITION 3.3. $\forall\, j \in OPT \setminus ON,\ v^*_{f(j)} \geq v(j)$.

Proof. Let $t = f(j)$. First notice that $j \in A_t$: $j \notin ON$, $r(j) \leq t$ as $j \in OPT[1,t]$, and $d(j) \geq t$ since either $j \in X_{t-1}$ or $j = OPT[t]$. Let $m_t \in S_t$ be the player who determined v^*_t, (if $v^*_t = 0$ then set $m_t = null$, so $S_t \setminus m_t = S_t$). We first show that, by subclaim 3.1, $j \notin S_t \setminus m_t$: define A as $OPT[1,t] \setminus ON$ minus all players with deadline $< t$, and $B = A_t$. Clearly $A \subseteq B$. By definition, X_t is optimal for A (over items $t+1, ..., M$). $S_t \setminus m_t$ is optimal for B (over items $t+1, ..., M$): if $m_t = null$ this follows from the optimality of S_t, and if $m_t \neq null$ this follows from sub-claim 3.2. Therefore, since $j \notin X_t$ then $j \notin S_t \setminus m_t$. If $j \neq m_t$ then $j \notin S_t$, and since $j \in A_t$ it follows from the optimality of S_t that $v(j) \leq v(m_t)$. If $j = m_t$ then this trivially holds. Therefore $v(j) \leq v(m_t) = v^*_{f(j)}$, and the claim follows.

PROPOSITION 3.4. *f is 2 to 1.*

Proof. Fix any time t. We need to show that f maps at most two players to t. Let $j_1 \in X_{t-1}$ be the player with minimal value such that $X_{t-1} \setminus j_1$ is an allocation of items $t+1, ..., M$, and denote $Y = X_{t-1} \setminus j_1$ (if X_{t-1} itself is independent w.r.t items $t+1, ..., M$ then set $Y = X_{t-1}$). If $X_t \subseteq Y$ then by the optimality of X_t it follows that $X_t = Y$ and the claim follows: by definition, f maps only j_1 and $OPT[t]$ to t. Otherwise, $X_t \setminus Y \neq \emptyset$. We first show that $X_t \setminus Y = \{OPT[t]\}$. This is implied by sub-claim 3.1: set $A = OPT[1, t-1] \setminus ON$, and $B = OPT[1,t] \setminus ON$. Since Y is optimal for A (by sub-claim 3.2) and X_t is optimal for B (by definition) it follows that, if $j \in OPT[1, t-1]$ but $j \notin Y$ then $j \notin X_t$, i.e. that $X_t \setminus Y = \{OPT[t]\}$. To conclude, we observe that X_t is a base in the matroid over items $t+1, ..., M$ and players $OPT[1,t] \setminus ON$, and that Y is an independent set of that matroid. Therefore $|Y \setminus X_t| \leq |X_t \setminus Y| = 1$, and thus $|X_{t-1} \setminus X_t| \leq 2$. Since $OPT[t] \in X_t$ then, by definition, the players mapped to t are exactly those in $|X_{t-1} \setminus X_t|$, and the claim follows.

4 The Impossibility of Truthful Approximations

We now move from algorithmic considerations to game-theoretic ones, in order to analyze player strategies. Since our goal is to find approximately optimal allocations with respect to the *true* variables of the players, we would prefer to design truthful algorithms: an algorithm is truthful if there exist price functions such that, regardless of how the other players act, player i will maximize his utility by declaring his true variables to the algorithm[10]. More formally, let T_i be the domain of all valid player i types/bids $(r(i), v(i), d(i))$, and let $T_{-i} = \times_{j \neq i} T_j$. Consider the allocation constructed by the algorithm upon receiving the type $b_i \in T_i$ from player i and $b_{-i} \in T_{-i}$ from the other players, and let $v(i, b)$ be the value that player i obtains from this allocation, i.e. $v(i)$ if i receives one of his desired items, and 0 otherwise.

DEFINITION 4.1. (TRUTHFULNESS) *An algorithm is truthful if there exist price functions $p_i : T_1 \times \cdots \times T_n \rightarrow \Re$ such that, for any i, any $b_{-i} \in T_{-i}$, any true type $b_i \in T_i$, and any $\tilde{b}_i \neq b_i$[11]:*

$$v(i, b_i, b_{-i}) - p_i(b_i, b_{-i}) \geq v(i, \tilde{b}_i, b_{-i}) - p_i(\tilde{b}_i, b_{-i}).$$

Such a property is highly desirable, as it guarantees that each player will be motivated to reveal his true type to the algorithm, by an argument similar to the traditional worst-case arguments of Computer Science. Indeed, many recent examples show truthful algorithms for various CS models (some are cited in our introduction). However, for our model, no such algorithm performs well:

THEOREM 4.1. *Any truthful deterministic algorithm for our online allocation problem cannot obtain an approximation ratio better than M.*

Proof. Assume w.l.o.g. that a player that does not win any item pays 0. This implies that i's price must not be higher than his value.

LEMMA 4.1. *Fix some truthful deterministic mechanism with some fixed approximation ratio. Then, for any player i with $r(i) = 1$ there exists a price function $p_i : T_{-i} \rightarrow \Re$ such that, for any combination of players that arrive at time 1, b_{-i}:*

- *If $v(i) > p_i(b_{-i})$ then i wins item 1 and pays $p_i(b_{-i})$ (regardless of his deadline).*

[10] By the revelation principle [24], it is w.l.o.g. to consider only algorithms that receive, as input, the players' types.

[11] We actually restrict the possible $\tilde{b}_i$'s such that $\tilde{r}_i \geq r_i$.

- *If $v(i) < p_i(b_{-i})$ then i does not win any item.*

Proof. Fix any combination of players that arrive at time 1, b_{-i}. Suppose first that i has deadline equal to 1. For this case, the player becomes one parameter, and by truthfulness there exist a price function according to the claim [1][12].

We now show that this function p_i satisfies the conditions of the claim, regardless of i's deadline. Fix any deadline $d(i)$ of i. If $v(i) > p_i(b_{-i})$ then i must win some item until his deadline, otherwise he can declare $\tilde{d}_i = 1$ and have strictly better utility. But then, if i does not win item 1, the adversary will produce players with higher and higher values, forcing the mechanism not to allocate any item to i in order to maintain the approximation ratio, thus contradicting truthfulness. Therefore i will receive item 1. He will pay $p_i(b_{-i})$ as otherwise, if he pays a higher price, he will declare $\tilde{d}_i = 1$ and will reduce his price, and if he pays less, then if i will have deadline equals 1 he will declare $d(i)$ instead, thus still winning item 1 but paying less. Therefore the function p_i satisfies the first condition.

Suppose now that $v(i) < p_i(b_{-i})$, and suppose there exists a scenario in which i wins one of his desired items. His price must be at most $v(i) < p_i(b_{-i})$. But then, if i had some value larger than $p_i(b_{-i})$ he would have been better off declaring $v(i)$ instead, by this still winning but paying less. Therefore i cannot win any item at all, and the claim follows.

We now finish the proof of the theorem. Fix any price functions $p_i : T_{-i} \to \Re$. For any $\epsilon > 0$ we show the existence of player types $b_1, ..., b_M$ s.t. $\forall\, i: \; r(i) = 1$, $d(i) = M$, $1 \le v(i) \le 1 + \epsilon$, and $v(i) \ne p_i(b_{-i})$. By the above claim, it follows that the mechanism can obtain welfare of at most $1 + \epsilon$, while the optimal allocation is at least M, and the theorem follows. To verify that such types exist, fix $L > M$ real values in $[1, 1 + \epsilon]$. Choose M values $v(i)$ uniformly at random from these L values. Then, for any given i, $Pr(v(i) = p_i(v(-i))) \le 1/L$, as the values were drawn i.i.d. Thus, $Pr(\exists i,\ v(i) = p_i(v(-i))) \le M/L < 1$, hence there exist a choice of values with $v(i) \ne p_i(v(-i))$ for all i.

Remark 1: Although the proof utilizes an extreme scenario with players with very large values, the worst case ratio occurs in common, simple scenarios, as the proof demonstrates. I.e., since the algorithm defends itself against such extremes, it must make wrong decisions even in simple cases.

[12] The argument essentially states that, if i wins for some $v(i)$ then he ins with any higher value, and pays the same. Therefore there exists a threshold value $p_i^* = p_i^*(b_{-i})$, such that i wins and pays p_i^* if $v(i) > p_i^*$, and looses otherwise.

Remark 2: A simple truthful deterministic M-approximation exists: For any player i, set p_i to be the highest bid received in time slots $1, ..., t$, excluding i's own bid. Sell item t to player i if and only if $v(i) > p_i$, for a price of p_i.

5 Will Players Act As Expected?

Our main motivation at this point is to justify the assumption that players will behave "as expected". We desire a rational justification, i.e. one that shows that expected strategies are, in some sense, utility maximizers for the players. The settings that we are interested in are ones in which "recommended" strategies are indeed to be intuitively expected, and deviating from them would seem to require some effort. In such cases, even rather weak notions of rational justification carry some weight. Such settings include, in particular, situations where computer protocols are announced and appropriate software that acts "as expected" is available. Our notions are intended for cases where the existing standard notions do not apply: truthfulness is impossible, and no distributional assumptions can be made since we seek "worst-case" notions as in computer science.

5.1 A Game Theoretic Framework: Implementation in Set-Nash Equilibria

Our setting contains n players, where each player i has a privately known type (input) $t_i \in T_i$, and $T = T_1 \times ... \times T_n$. Each player i has a strategy space S_i, and a utility function $u_i : T_i \times S \to \Re$, where $u_i(t_i, s_i, s_{-i})$ denotes i's payoff when his type is t_i, he plays strategy s_i and the others play the strategy tuple s_{-i}. We model a situation in which a set of recommended strategies is defined for each player. Specifically, a function $R_i : T_i \to 2^{S_i}$ is given, where $R_i(t_i) \subseteq S_i$ is the set of strategies that player i may be expected to follow. We denote also $R_i(*) = \cup_{t_i \in T_i} R_i(t_i)$. The motivating scenario is where it is known that if all players i play strategies $s_i \in R_i(t_i)$, then the outcome is "good" in some sense. E.g., in our case, the obtained social welfare approximates the optimal one. We would like to capture the notion that the sets R_i are in equilibrium. In other words, formalize when can it be said that given that other players $j \ne i$ all play strategies in $R_j(t_j)$, then player i also rationally plays some strategy in $R_i(t_i)$.

Definition 5.1. *The set functions $\{R_i(\cdot)\}$ are in* **Set-Nash equilibrium** *(for pure strategies) if for every i, every $t_i \in T_i$, every $s_{-i} \in R_{-i}(*)$, and every $s_i \in S_i$ there exists $r_i \in R_i(t_i)$ such that $u_i(t_i, r_i, s_{-i}) \ge u_i(t_i, s_i, s_{-i})$. I.e. if all others play some recommended strategies (not necessarily according to their true types) then there exists a best response strategy for i that is* *one of i's recommended strategies (according to his true type).*

Although this definition is weak, we believe that this "set equilibrium" concept is important. Specifically, it captures the intuitive rational behind playing a "semi-myopic" strategy in our model, as we show below. To answer the weak points of the definition, while still maintaining its spirit, we suggest a hierarchy of strengthened definitions in the full paper, and discuss their different properties.

It seems most appropriate to formalize our main theorem using the framework of Implementation Theory. This will also help us to describe the structure of our proof, below. In this setting, we have a set of outcomes/alternatives, A, from which we have to choose one outcome. The choice depends on the players types $t \in T$, according to some social rule $F : T \rightarrow 2^A$. In our case, A is the set of all valid allocations of items to players, and $F(t)$ outputs all allocations that are 3-approximations w.r.t t. This social rule represents the fact that our goal is to obtain a 3-approximation of the welfare, and any allocation that obtains this will satisfy us. All the classic definitions from implementation theory can be adapted to our Set-Nash definition:

DEFINITION 5.2. *Given $F : T \rightarrow 2^A$, an implementation in Set-Nash equilibrium is a mechanism with strategy sets $S_1, ..., S_n$, and an outcome function $g(s_1, ..., s_n) \in A$, such that there exists a Set-Nash equilibrium $\{R_i(\cdot)\}_i$ that satisfies that $g(s) \in F(t)$ for* all *$s \in R(t)$.*

The celebrated revelation principle states that whenever we can implement a social function in some equilibrium, we can also implement it using a direct revelation implementation, in which the strategy space of the players is simply to reveal their type. For our "set equilibrium" notion, we can have an "extended direct revelation" implementation which is "extended truthful":

DEFINITION 5.3. *An implementation is an "extended direct revelation implementation" if the strategies of the players are of the form (t_i, l_i), where $t_i \in T_i$, and l_i represents any additional information.*

An extended direct revelation implementation is "extended truthful" (in Set-Nash equilibrium) if there exists a Set-Nash equilibrium in which $R_i(t_i) = (t_i, *)$, i.e. the player declares his true type in every one of his recommended strategies.

PROPOSITION 5.1. *(An extended revelation principle) Every function $F : T \rightarrow 2^A$ that can be implemented in Set-Nash equilibrium can be implemented by an extended truthful implementation.*

5.2 Main Theorem The above Definitions enable us to exactly state the game-theoretic properties of our two ascending auctions:

Main Theorem: The Online Iterative Auction and the Sequential Japanese auction both implement a 3-approximation of the welfare in Set-Nash equilibrium.

In other words, both auctions have a Set-Nash equilibrium with the property that, for any combination of strategies from the recommended sets, the resulting welfare is at most one third of the optimal welfare.

We prove this in two steps. We first explicitly describe an extended direct revelation mechanism that implements a 3-approximation of the welfare. We then show that the this mechanism is "embedded" in both our auctions (which are not direct revelation mechanisms, of-course, and are seemingly different). We next give a short exposition to both these arguments. Exact definitions and proofs are given in the full paper [21].

The extended direct revelation mechanism, which we call a "semi-myopic mechanism", is as follows. Each player declares his type (arrival time, value, and deadline) plus an additional "false" deadline. At each time t, the mechanism computes the sets A_t, S_t, and f_t, which are the the natural parallels of the notions in definition 3.8, where the deadline of each player is taken to be his "false" deadline if that has not already passed, and his "true" deadline otherwise. The winner is chosen to be some player from f_t (this is actually a family of mechanisms, as the exact rule of choosing from f_t is not specified). We associate a "time-t-price" to every player in f_t, which equals the value of the highest player remained outside of S_t. The winner pays his maximal "time-t-price" over all time periods. The recommend strategy set $R_i(\cdot)$ of player i is to declare his true type plus any "false" deadline not higher than his true deadline. We have:

Lemma: For any player i, and any $s_{-i} \in R_{-i}(*)$, i has a best response to s_{-i} in $R_i(t_i)$.

Lemma: For any combination of recommended strategies, the semi-myopic mechanism is a semi-myopic algorithm.

Corollary: The semi-myopic mechanism Set-Nash implements a 3-approximation of the welfare.

The next step is to use this building block to prove the main claim (the argument below is repeated for each auction separately, in a similar way). We first show that the ascending auction is an "extension" of a semi myopic mechanism: Define a subset of the strategies of the ascending auction as follows. Each player chooses a false deadline, plays myopically with it until it expires, and then plays myopically with the true deadline. This creates an obvious mapping between the two strategy (sub)sets. We show that when players play such strategies in the ascending auction the result (allocation plus payments) satisfies the requirements of a semi myopic mechanism (this follows from the offline properties of the auctions). Thus we conclude that the desirable Set-Nash equilibrium exists in the ascending auction, when strategies are restricted in this manner. The last argument then shows that other strategies do not improve the situation of the player, i.e. do not contain strictly better actions. This concludes the proof.

Acknowledgements

We wish to warmly thank Yair Bartal and Ahuva Mu'alem for their help in the early stages of this work. We also thank Moshe Babaioff, Liad Blumrosen, Rica Gonen, Daniel Lehmann, and Motty Perry for many helpful discussions and comments.

References

[1] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42st Annual Symposium on Foundations of Computer Science*, 2001.

[2] B. Awerbuch, Y. Azar, and A. Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proc. of the 35th ACM Symposium on Theory of Computing*, 2003.

[3] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *Proc. of the 13th Symposium on Discrete Algorithms*, 2002.

[4] Y. Bartal, F. Y. L. Chin, M. Chrobak, S. P. Y. Fung, W. Jawor, R. Lavi, J. Sgall, and T. Tichy. Online competitive algorithms for maximizing weighted throughput of unit jobs. In *Proc. of the 21st Symposium on Theoretical Aspects of Computer Science*, 2004.

[5] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi-unit combinatorial auctions. In *Proc. of the ninth Conference on Theoretical Aspects of Rationality and Knowledge*, 2003.

[6] Y. Bartal and R. Lavi. Analysis of the β-opt algorithm, 2002. Unpublished manuscript.

[7] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. In *Proc. of the 14th Symposium on Discrete Algorithms*, 2003.

[8] A. Blum, T. Sandholm, and M. Zinkevich. Online algorithms for market clearing. In *Proc. of the 13th Symposium on Discrete Algorithms*, 2002.

[9] L. Blumrosen and N. Nisan. Auctions with severely bounded communication. In *Proc. of the 43rd Annual Symposium on Foundations of Computer Science*, 2002.

[10] G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.

[11] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63(1), 2001.

[12] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive generalized auctions. In *Proc. of the 34th ACM Symposium on Theory of Computing*, 2002.

[13] E. Friedman and D. Parkes. Pricing wifi at starbucks: issues in online mechanism design. In *Proc. of the ACM Conference on Electronic Commerce*, 2003.

[14] F. Gul and E. Stacchetti. The english auction with differentiated commodities. *Journal of Economic Theory*, 92:66–95, 2000.

[15] B. Hajek. On the competitiveness of on-line scheduling of unit-length packets with hard deadlines in slotted time. In *Proc. of the 2001 Conference on Information Sciences and Systems*, 2001.

[16] M. T. Hajiaghayi, R. Kleinberg, and D. Parkes. Adaptive limited-supply online auctions. In *Proc. of the ACM Conference on Electronic Commerce*, 2004.

[17] E. Horowitz and S. Sahni. *Fundamentals of Computer Algorithms*. Computer Science press, 1978.

[18] A. Kesselman, Z. Lotker, Y. Mansour, B. Patt-Shamir, B. Schieber, and M. Sviridenko. Buffer overflow management in qos switches. In *Proc. of the 33rd ACM Symposium on Theory of Computing*, 2001.

[19] R. Kleinberg and T. Leighton. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *Proc. of the 44rd Annual Symposium on Foundations of Computer Science*, 2003.

[20] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science*, 310:159–180, 2004.

[21] R. Lavi and N. Nisan. Online ascending auctions for gradually expiring items, 2004. http://www.cs.huji.ac.il/~tron/papers/online-ascending.ps.

[22] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proc. of the ACM Conference on Electronic Commerce*, 2001.

[23] D. Lehmann, L. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

[24] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford university press, 1995.

[25] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

[26] R. Porter. Mechanism design for online real-time scheduling. In *Proc. of the ACM Conference on Electronic Commerce*, 2004.

[27] M. P. Wellman, W. E. Walsh, P. R. Wurman, and J. K. MacKie-Mason. Auction protocols for decentralized scheduling. *Games and Economic Behavior*, 35:271–303, 2001.