An Online Multi-unit Auction with Improved Competitive Ratio

Sourav Chakraborty[1] and Nikhil Devanur[2]

[1] CWI, Amsterdam, Netherlands
sourav.chakraborty@cwi.nl
[2] Microsoft Research, Redmond.
nikdev@microsoft.com⋆⋆⋆

Abstract. We improve the best known competitive ratio (from 1/4 to 1/2), for the online multi-unit allocation problem, where the objective is to maximize the single-price revenue. Moreover, the competitive ratio of our algorithm tends to 1, as the bid-profile tends to "smoothen". This algorithm is used as a subroutine in designing truthful auctions for the same setting: the allocation has to be done online, while the payments can be decided at the end of the day. Earlier, a reduction from the auction design problem to the allocation problem was known only for the unit-demand case. We give a reduction for the general case when the bidders have decreasing marginal utilities. The problem is inspired by sponsored search auctions.

1 Introduction

It is fairly common that a mechanism has to work in a dynamic environment, where there is an uncertainity in either the demand, or the supply, or both. This has led to the study of *online mechanism design* [7, 4, 13] and has presented significant new challenges compared to the traditional static setting. Most of the research has focused on dynamic demand case: the uncertainty is in the number and types of the bidders, their arrival and departure time, etc, such as airline tickets. On the other hand, very little is known for dynamic supply case: the uncertainty is in the number of items to be allocated, or more generally the set of feasible allocations, such as sponsored search. Mahdian and Saberi [14] initiated the study of the dynamic supply case by giving a constant competitive ratio algorithm for auctioning multiple copies of a single item with unit-demand bidders.

⋆⋆⋆ Part of the work was done when both the authors were visiting Microsoft Research, India in Summer 2007 and when the first author was a graduate student at the University of Chicago and the second author was in Toyota Technological Institute, Chicago

We improve their competitive ratio by a factor of 2 by giving an alternate and simple algorithm, and also extend their results to handle bidders with multiple demand.

A bidder with *unit-demand* has a value u_i for one copy of the item, and his utility is $u_i - p$ if he is allocated the item at price p, and 0 otherwise.

Definition 1. Online Multi-unit Auction Problem, unit-demand *([14]). At the beginning of the auction, each bidder (with unit-demand) bids a value b_i. At each (discrete) time unit, a new copy arrives which must be allocated to a bidder immediately, or else it perishes. When there are no more copies left, the auction determines the prices charged to the winning bidders.*

Note that the auction has no prior knowledge of how many copies of the item will be produced.

Definition 2. *An auction is* truthful *if bidding $b_i = u_i$ is a dominant strategy for each bidder.*

The goal of the auction is to maximize the revenue of the auctioneer, which is the sum of the prices charged to the winning bidders. The main motivation behind the work of [14] was sponsored search auctions, which are a major source of revenue for search engines like Google, Yahoo and MSN. The bidders are the advertisers and the items correspond to search queries. The queries arrive online and have to be allocated immediately, while the advertisers stay for the entire duration of the auction and present their bids ahead of time. The advertisers are only charged at the end of the day.

An alternate model is to ask that the prices are also determined online. For the sponsored search auction setting, charging at the end is closer to reality. Also, charging online seems to be considerably restrictive, as there are strong lower bounds for this model[1]. For the sponsored search auction setting, a more realistic model is when the bidders have multiple demand. We present an auction for this case as well, and our results for this case are of significant interest.

The auction problem considered here is also a natural extension of the line of work on digital goods auction: from unlimited supply ([9–12]) to limited supply ([1, 3, 5]), to unknown supply ([14] and this paper).

[1] The lower bounds [2] are for a related problem, that of maximizing social welfare. It is an interesting open question if these lower bounds also hold for maximizing revenue.

As is standard in the literature on digital goods auction, we give a competitive analysis of the auction, by comparing the revenue of the auction to a benchmark. The benchmark we use is once again a standard in digital goods auction, it is the *optimal single-price revenue* on hindsight: $OPT := \max_p p.|\{i : b_i \geq p\}|$. The auction itself is allowed to charge different prices to different bidders, although our auction charges only two different prices.

Definition 3. Competitive Ratio *An auction is said to have a competitive ratio of α if the expected revenue of the auction is at least αOPT.*

[14] gave a reduction from the auction problem to the following algorithmic problem, with only a constant factor lost in the competitive ratio.

Definition 4. Online Multi-unit Allocation Problem, unit-demand *The algorithm is given the utility u_i of each bidder. At each (discrete) time unit, a new copy arrives which must be allocated to a bidder immediately, or else it perishes. When there are no more copies left, the algorithm charges all the winning bidders with a* single *price, that is smaller than their utilities.*

There is no requirement of truthfulness in the allocation problem. Also, the algorithm itself has to charge the same price to all the bidders, unlike the auction which was allowed to charge different prices to different bidders. Also note that the revenue-maximizing single price is determined by the allocations made by the algorithm. It is simply the smallest winning utility. As with the auction problem, we compare the revenue of the algorithm with the optimal single price revenue on hindsight; the competitive ratio is defined analogously.

Theorem 1. *([14]) There is a truthful mechanism for the online multi-unit auction problem with unit-demand bidders with a competitive ratio of $O(\alpha)$ given an algorithm for the allocation problem with competitive ratio α.*

1.1 Main Result

It can be easily seen that the competitive ratio of any deterministic algorithm for the allocation problem is arbitrarily small. So it is actually surprising that a randomized algorithm can even get a constant competitive ratio. The reason for this difficulty is that the revenue of the algorithm, as a function of the number of copies allocated can have many "peaks" and "valleys". For any deterministic algorithm, an adversary can make sure that the algorithm either ends up in a valley, or is stuck on a small peak while the optimum is at a larger peak elsewhere. The key decision for an algorithm is when it is at a peak, it has to decide if it has to stay at the peak, or try to get to the next one. What our algorithm does is to simply wait at the current peak for a period of time chosen uniformly at random between 1 and the maximum distance between peaks seen so far. The simplicity of our algorithm is quite appealing. This improves the best known competitive ratio (from 1/4 to 1/2), for the online multi-unit allocation problem, which in turn gives a factor of 2 improvement for the online multi-unit auction problem.

Theorem 2. *There is an algorithm for the online multi-unit allocation problem for unit demand, that achieves a competitive ration of* $1/2$*.*

The proof of the competitive ratio relies on case analysis since the optimal revenue and the expected revenue of the algorithm vary depending on the total number of copies seen. A good idea of how the analysis goes can be had by considering the following instance: suppose there is one bid of 1 and many bids of $\epsilon \ll 1$. In this case the algorithm waits for a time chosen u.a.r between 1 and $1/\epsilon$. If the number of copies seen is $m \leq 1/\epsilon$, then the optimal revenue is 1, while the expected revenue is $1 - x + \frac{x^2}{2}$ (where $x = \epsilon m$), which is at least $1/2$ when $x \leq 1$. If $m \geq 1/\epsilon$ then the optimal revenue is ϵm, while the expected revenue is $\epsilon m - 1/2 \geq \frac{\epsilon m}{2}$.

Moreover, the competitive ratio of our algorithm tends to 1, as the bid-profile tends to "smoothen". [14] also showed an upper bound of $e/(e+1)$ for the allocation problem and closing the gap is an open problem. See Section 4 for a more detailed discussion on this.

Mahdian and Saberi [14] showed that using an algorithm for the online multi-unit allocation problem for unit demand with competitive ratio ρ one can construct a truthful auction for the online multi-unit auction problem with competitive ratio $\rho/20$. Thus,

Corollary 1. *There is a truthful auction for the online multi-unit auction problem, that achieves a constant competitive ratio.*

1.2 Extensions

A more realistic case in the context of sponsored search auction is when the bidders have multiple demand: bidders have decreasing marginal utilities for multiple copies of the item, and submit multiple bids. The optimum and the competitive ratio are defined analogous to the unit-demand case. The allocation problem remains the same even with multiple demands, since the problem does not really depend on the identity of the bidders. Hence, our algorithm for the allocation problem gives a competitive ratio of 1/2 even for this case.

However, the auction problem is harder with multiple demands, since it provides more ways for the bidders to lie and benefit. In particular, the auction obtained by using the reduction in [14] is not truthful for multiple demands. The reduction in [14] is based on random sampling with computing optimal "price offers". But when run in an online setting, the prices offered decrease over time, due to which a bidder might regret not getting a copy earlier as the price decreased at a later time. The reduction in [14] takes care of this situation by a clever implementation that works only when all bidders want only one copy. It is not truthful when the bidders can submit multiple bids. We circumvent this difficulty by combining the random sampling technique with the VCG auction. However, we only get an asymptotic competitive ratio, that is the ratio tends to 1/2, as a certain bidder dominance parameter tends to 0. The bidder dominance parameter is defined to be the maximum fraction of the optimum revenue that can be obtained from any single bidder. A small bidder dominance parameter indicates that the revenue from any one bidder is small compared to the optimal revenue.

Definition 5. *For any price p and any bidder i we denote by $n(i,p)$ the number of bids of bidder i that are more than p. The bidder dominance parameter is*

$$\eta := \frac{\max_{i,p} n(i,p)p}{OPT}.$$

Theorem 3. *There is a truthful mechanism for the online multi-unit auction problem with multiple-demand bidders, that with probability more than $(1-\delta)$ guarantees a revenue of at least $\alpha OPT(1-\epsilon)$ on expectation, where α is the competitive ratio of the allocation algorithm that we use as the subroutine, if*

$$\eta = O\left(\epsilon^2 / \log\left(\frac{n}{\delta}\right)\right),$$

where n is the number of distinct bid values.

The problem considered here is perhaps the simplest non-trivial case of the actual problem in sponsored search auctions. There are many extensions of which we have little understanding, for instance, one could consider multiple slots for every query. Another interesting extension is when the bidders have constant marginal utilities for the copies, but have daily budgets. [5, 1] gave an auction for this case with known supply (the offline problem). Extending it to the online setting is an important open problem. The introduction of budgets also makes the multiple items case interesting. (Otherwise, assuming additive utilities, the auctions for different items are independent of each other.) Even the offline case of this problem is open.

Subsequent Related Work: Subsequent to our result, Devanur and Hartline [6] gave an alternate auction for the Online Multi-Unit Auction problem with a competitive ratio that is better than this paper. This auction does not use the reduction to the allocation problem. However, the auction in [6] is only for unit-demand bidders, so our results for the multiple-demand bidders are still the best. Also the online allocation problem, and the algorithm for it are interesting in their own right.

Organization: We present our algorithm for the Online Multi-unit Allocation problem in Section 2. Theorem 2. For lack of space we are unable to present the proof of the competitive ratio of the algorithm in this extended abstract. The auction for the multiple demands case and a sketch of the proof of Theorem 3 is given in Section 3. Section 4 contains a discussion on future work and open problems.

2 Algorithm for the Online Multi-unit Allocation Problem

Without loss of generality assuming that the utilities are $u_1 \geq u_2 \geq \cdots \geq u_n$, the revenue obtained by allocating l units of the item is lu_l. Let $1 = a_1 < b_1 < a_2 < b_2 < a_3 < b_3 < \ldots$ be the critical points of the function lu_l, that is, the function lu_l is non-decreasing as l increases from a_i to b_i and for all $b_i < l < a_{i+1}$ we have $b_i u_{b_i} > lu_l$ and $b_i u_{b_i} \leq a_{i+1} u_{a_{i+1}}$.

The algorithm is in one of two states, ALLOCATE or WAIT. When it is in ALLOCATE, it allocates the next copy of the item. When it is in WAIT, it discards the next copy. The description of the algorithm is completed by specifying when it transits from one state to the other.

The algorithm is initially in ALLOCATE. It transits from ALLOCATE to WAIT when the number of copies allocated (X) is equal to b_i for some i. It transits from WAIT to ALLOCATE when the number of copies discarded till then (Y) is equal to a random variable, T, for waiting time. T is reset every time the algorithm transits to WAIT. T is picked so that it is distributed uniformly between 0 and D_i, where $D_0 = 0$ and for all $i \geq 1$

$$D_i = \max_{j \leq i}(a_{j+1} - b_j)$$

(recall that $X = b_i$). We further want to maintain the invariant that Y never exceeds T. Equivalently, the value of T can only increase during a run of the algorithm.

We still have to specify how T is picked. Because of the condition that T can only increase, we cannot pick T independently every time we transit to WAIT. If $D_i \leq D_{i-1}$, then we don't have to change T at all. If $D_i > D_{i-1}$, then

- w.p. $\frac{D_{i-1}}{D_i}$ don't change T,
- with the remaining probability pick T uniformly at random from the interval $[D_{i-1}, D_i]$.

It is easy to see that the resulting T is distributed uniformly in $[0, D_i]$. Note that in case T is not changed, then Y is already equal to T, and we transit back to ALLOCATE immediately. Equivalently, we don't transit to WAIT at all.

Pseudocode for the Algorithm

```
initialize STATE = ALLOCATE, i=1, X=Y=T=0;
when a new copy is produced
    If (STATE = ALLOCATE)
        Allocate the copy to the next bidder;
        X ++;
        If (X = b_i)
            If (D_i > D_{i-1})
                With prob 1 - D_{i-1}/D_i
                    set T to a random number from the interval [D_{i-1}, D_i];
                    STATE = WAIT;
            i ++;
    If (STATE = WAIT)
```

```
13.            Discard the copy;
14.            Y ++;
15.            If (Y = T)
16.                STATE = ALLOCATE
17.        GO TO line 2.
```

Because of shortage of space we cannot present the analysis for competitive ratio in this extended abstract.

3 Bidders with Multiple Demand

Let $\mathbf{B} = \{1, 2, \ldots, n\}$ be the set of bidders. Each bidder can make multiple bids. We will design a truthful mechanism which has good competitive ratio. Our mechanism will use an online multi-unit allocation algorithm as a sub-routine. Under a bidder-dominance assumption, the competitive ratio of our mechanism will be $(1 - \epsilon)\alpha$ where α is the competitive ratio of the allocation algorithm we use as our subroutine.

The Mechanism: We divide the set of bidders into two groups S and T by placing each bidder randomly into either of the groups. On each set of bidders S and T we will have fictitious runs of the allocation algorithm. Let the fictitious run of the allocation algorithm on the set S (respectively T) allocates $x(S, k)$ (respectively $x(T, k)$) copies when k copies are produced.

Now when the j-th copy is produced, if j is even we compute $x(S, j/2)$. If at that time the number of copies allocated to bidders in T is less than $x(S, j/2)(1 - 6\gamma)$ then we allocate the j-th copy to T otherwise discard the copy. Similarly, if j is odd we compute $x(T, (j+1)/2)$ and if the number of copies allocated to bidders in S is less than $x(T, (j + 1)/2)(1 - 6\gamma)$ then we allocate the j-th copy to S otherwise discard the copy.

Finally let $x_{final}(S)$ and $x_{final}(T)$ copies are allocated to bidders in S and T respectively. The prices charged are the VCG payments, that is, as if we ran a VCG auction to sell $x_{final}(S)$ copies to bidders in S.

Note that the even indexed copies will be allocated only to bidders in T and the odd-indexed copies will be allocated only to bidders in S. But the bids of bidders in S decides how many (odd-indexed) copies will be allocated to bidders in T and vice versa. This mechanism is similar to that in [11] on digital good auction with unlimited supplies except that in [11] the bids of bidders in S decides the cut off price for bidders in T and vice-versa.

If M is the number of copies of the item that are finally produced we denote by $OPT = OPT(\mathbf{B}, M)$ the revenue obtained by the optimal single price allocation algorithm.

Definition 6. *For any price p and any bidder i we denote by $n(i,p)$ the number of bids of bidder i that are more than p.*

We define the bidder dominance parameter η as

$$\eta = \frac{\max_{i,p} n(i,p)p}{OPT}.$$

Theorem 4. *The above mechanism is a truthful mechanism. If all the bids are from a finite set of prices (say Q) and if*

$$\frac{1}{\eta} = \Omega\left(\log\left(\frac{|Q|}{\delta}\right)\left(\frac{1}{\epsilon^2}\right)\right)$$

and if we set $\gamma = \epsilon/8$ then with probability more than $(1-\delta)$ our mechanism guarantees a revenue of at least $\alpha OPT(1-\epsilon)$ on expectation, where α is the competitive ratio of the allocation algorithm that we use as the subroutine.

In the rest of this section we will give a sketch of the proof of the theorem. The detailed proof of the theorem is in the Appendix. The proof is similar to that in [11].

The proof that the mechanism is truthful follows from the facts that the number of copies allocated to each half is independent of the number of the bids of the bidders in that half and the fact that pricing is determined by the VCG auction.

The proof of the competitive ratio has two main parts: The first thing is that since the bidders are split randomly into two sets so with high probability the optimal revenue we can obtain from either of the sets is nearly half of what we can obtain from the whole set.

The second thing is that the discounting factor of $(1-6\gamma)$ ensures that with high probability the eventual winners in S (respectively T) are charged at least as much as our allocation algorithm charges during its fictitious run on the set T (respectively S) .

Note that the bound on the bidder dominance gives us an upper bound on $n(i,p)$ that is the number of bids on any bidders that is more than p. This is essential for our analysis.

Let a fictitious run of the optimal single price allocation algorithm on S generates a revenue of $OPT(S,j)$ after j copies are produced. By Mcdiarmid's Inequality and the bound on the bidder dominance parameter, with probability at least $(1-O(\delta))$ we have $OPT(S,\lceil M/2\rceil) > (1/2-\gamma)OPT$, where M is the final number of copies produced. Similarly we have $OPT(T,\lfloor M/2\rfloor) > (1/2-\gamma)OPT$.

For the second stage we again notice that since the set of bidders was partitioned randomly so with high probability the set of bids that are more than p is also evenly divided among the two sets S and T. From the McDiarmid's Inequality and from the bound on the bidder dominance parameter we see that with high probability the number of bids in S that are more than p is much more than $(1-6\gamma)$ times the number of bids in T that are more than p (and vice versa).

Let $ALG(S,j)$ and $ALG(T,j)$ be the revenue is generated by the fictitious run of our allocation algorithm on S and T respectively after j items are produced. Now since the allocation algorithm is α competitive we have that on expectation $ALG(S,j) > \alpha OPT(S,j)$. Thus with probability at least $(1-O(\delta))$ the revenue we earned on expectation is more than

$$ALG(S,\lceil M/2\rceil)(1-6\gamma)+ALG(T,\lfloor M/2\rfloor)(1-6\gamma) > \alpha(1-6\gamma)(1-2\gamma)OPT$$

which is greater than $\alpha(1-8\gamma)OPT$.

4 Conclusion and Open Problems

The optimal competitive ratio for the allocation problem is open. [14] showed an upper bound of $e/(e+1)$ for any randomized algorithm. The instance for which they show this upper bound is when there is one bid of 1 and many bids of ϵ. For this particular instance, the following algorithm gets a competitive ratio of 2/3: with probability 1/3, allocate just one copy and get a revenue of 1, and with probability 2/3, run our algorithm. We conjecture that this algorithm can be generalized to get a 2/3 competitive ratio. Also, a better upper bound proof will probably have to consider instances with multiple peaks, where the ratio of the D_i's to the a_{i+1}'s is large.

For the auction problem, the competitive ratio for the unit-demand case is quite small, and that for the multiple demand case holds only asymptotically. Getting it to a reasonably large constant (or proving that it is impossible) is an important open problem.

The most common scenario in sponsored search auctions is that the bidders have a constant utility for multiple copies of the item, but with a daily budget. Our allocation algorithm works for this case as well, but the reduction from the auction problem is not truthful. Borgs et al [5] give a truthful auction for the offline case with budgets, using the standard random sampling techniques with price offers. However, it is not clear how to extend their auction to the online case. The difficulty is the same as that for the multiple-demand case, that the price offers are decreasing over time. But unlike the multiple-demand case, there is no VCG auction for the budgets case, so our reduction does not work.

5 Acknowledgements

The authors would like to thank Jason Hartline for useful discussions about the digital goods auction, and Deeparnab Chakrabarty for comments on an earlier draft of the paper.

References

1. Zo Abrams: *Revenue maximization when bidders have budgets.* SODA 2006: p.1074-1082
2. Moshe Babaioff, Liad Blumrosen, Aaron L. Roth: *Auctions with Online Supply* Manuscript.
3. Maria F. Balcan, Nikhil R. Devanur, Jason Hartline, Kunal Talwar: *Random Sampling Auctions for Limited Supply.* Manuscript, submitted.
4. Avrim Blum, Jason D. Hartline: *Near-optimal online auctions.* SODA 2005: p.1156-1163
5. Christian Borgs, Jennifer T. Chayes, Nicole Immorlica, Mohammad Mahdian, Amin Saberi: *Multi-unit auctions with budget-constrained bidders.* ACM Conference on Electronic Commerce 2005: p.44-51
6. Nikhil R. Devanur, Jason Hartline: *Limited and Online Supply and the Bayesian foundations of prior-free mechanism design.* To appear in ACM Conference on Electronic Commerce 2009.
7. David Parkes, *Online mechanisms.* In Algorithmic Game Theory, Noam Nisan, Tim Roughgarden, Eva Tardos, Vijay Vazirani (eds.) Chapter 16. CUP, 2007.

8. Uriel Feige, Abraham Flaxman, Jason D. Hartline, Robert D. Kleinberg: *On the Competitive Ratio of the Random Sampling Auction*. WINE 2005: p.878-886
9. Andrew V. Goldberg, Jason D. Hartline: *Envy-free auctions for digital goods.* ACM Conference on Electronic Commerce 2003: p.29-35
10. Andrew V. Goldberg, Jason D. Hartline: *Competitiveness via consensus.* SODA 2003: p.215-222
11. Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Michael E. Saks, Andrew Wright: *Competitive auctions.* Games and Economic Behavior (55) 2006: p.242-269.
12. Andrew V. Goldberg, Jason D. Hartline, Andrew Wright: *Competitive auctions and digital goods.* SODA 2001: p.735-744
13. Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, David C. Parkes: *Adaptive limited-supply online auctions.* ACM Conference on Electronic Commerce 2004: p.71-80
14. Mohammad Mahdian, Amin Saberi: *Multi-unit Auction with Unknown Supply.* ACM Conference on Electronic Commerce 2006: p.243-249.

Appendix

A Competitive Analysis of the Algorithm for the Online Multi-unit Allocation Problem

In this section we show that the expected revenue of our algorithm, ALG, is at least half of the optimal revenue on hindsight, OPT. Let M be the number of copies that is produced at the end of the day. If $M \leq b_1$ then OPT is Mb_M and our algorithm also allocates all the M copies. Thus in this case ALG achieves the optimal revenue. Now let $b_i \leq M < b_{i+1}$.

Case 1: If $M \leq a_{i+1}$, then OPT is $b_i u_{b_i}$.
Case 2: If $a_{i+1} < M$, then OPT is Mu_M.

Recall that X is the number of items sold by the algorithm. Therefore, $ALG = E[Xu_X]$. We approximate ALG in the above two cases as follows. In Case 1, if $X \leq b_i$ then $u_X \geq u_{b_i}$, and if $X > b_i$ then $u_X \geq u_{a_{i+1}}$. Hence we have that

$$ALG \geq \Pr[X < b_i]E[X|X < b_i]u_{b_i} + \Pr[X = b_i]b_i u_{b_i} + \Pr[X > b_i]E[X|X > b_i]u_{a_{i+1}}$$

Also note that by definition $a_{i+1}u_{a_{i+1}} \geq b_i u_{b_i} = OPT$. Hence

$$u_{a_{i+1}} \geq \frac{b_i u_{b_i}}{a_{i+1}} = \frac{OPT}{a_{i+1}}.$$

Substituting for the values of u_{b_i} and $u_{a_{i+1}}$, we get

$$\frac{ALG}{OPT} \geq \Pr[X < b_i]E[X|X < b_i]\frac{1}{b_i} + \Pr[X = b_i] + \Pr[X > b_i]E[X|X > b_i]\frac{1}{a_{i+1}}. \quad (1)$$

In Case 2, we use the fact that since $X \leq M$ thus $u_X \geq u_M$, so $ALG \geq E[X]u_M$, and since $OPT = Mu_M$, we need to prove that $E[X]$ is at least $M/2$.

$$E[X] = \Pr[X < b_i]E[X|X < b_i] + \Pr[X = b_i]b_i + \Pr[X > b_i]E[X|X > b_i]. \quad (2)$$

We now give a way to calculate the various probabilities and expectations needed.

Definition 7. *For all $i \geq 1$ let T_i be the value of the random variable T chosen at phase i.*

That is, T_i is the number of number of items we plan to discard before allocating the $(b_i + 1)$-th element. Also note that T_i is distributed uniformly between 0 and D_i, and $T_{i-1} \leq T_i$ for all i. Also, it is easy to see from the description of the algorithm that $X \geq M - T_{i-1}$ when $X < b_i$, and $X = M - T_i$ when $X > b_i$. This gives us the following lemmas. Let $M' := M - b_i$.

Lemma 1. *For all $i \geq 1$ the following statements hold, (let $T_0 = 0$),*

1. $X < b_i \Leftrightarrow T_{i-1} > M'$.
2. $X = b_i \Leftrightarrow T_{i-1} \leq M'$ *and* $T_i \geq M'$.
3. $X > b_i \Leftrightarrow (T_{i-1} < M'$ *and* $T_i < M') \Leftrightarrow T_i < M'$.

Lemma 2.

$$E[X|X < b_i] \geq M - E[T_{i-1}|T_{i-1} > M'] = M - \frac{M' + D_{i-1}}{2}.$$
$$E[X|X > b_i] = M - E[T_i|T_i < M'] = M - \frac{\min\{M', D_i\}}{2}.$$

However, the probability of the events $X < b_i$, $X = b_i$ and $X > b_i$ depend upon the order of $J_i := b_i + D_{i-1}$, a_{i+1} and M. So we consider all possible orders of these 3 quantities separately. Table 1 shows the probabilities for all the cases.

A.1 Analyzing all the Cases

A few observations first: from Lemma 2,

$$E[X|X < b_i] \geq M - \frac{M' + D_{i-1}}{2} \geq \frac{M}{2}$$

since $D_{i-1} \leq b_i = M - M'$ and hence $M' + D_{i-1} \leq M$. And when $M' \leq D_i$,

$$E[X|X > b_i] = M - \frac{M'}{2} = \frac{M + b_i}{2}.$$

Case 1a: $[J_i \leq M \leq a_{i+1}]$ Set $x = \frac{M'}{D_i}$ and $y = \frac{D_i}{a_{i+1}}$. Then $\frac{M'}{a_{i+1}} = xy$ and $\frac{b_i}{a_{i+1}} = 1 - y$. Substituting for the probabilities and expectations in (1),

$$\frac{ALG}{OPT} \geq 1 - x + x\left(1 + \frac{xy}{2} - y\right) = 1 + \frac{x^2 y}{2} - xy =: \alpha.$$

$$\frac{d\alpha}{dx} = y(x - 1) \leq 0.$$

Therefore α is minimized when $x = 1$, and at this point, $\alpha = 1 - y/2 \geq 1/2$ since $y \leq 1$.

	$\Pr[X < b_i]$ $= \Pr[T_{i-1} > M']$	$\Pr[X = b_i]$	$\Pr[X > b_i]$ $= \Pr[T_i < M']$
Case 1a $J_i < M \leq a_{i+1}$ $D_{i-1} \leq M' \leq D_i$	0	$1 - \frac{M'}{D_i}$	$\frac{M'}{D_i}$
Case 1b $M \leq J_i < a_{i+1}$ $M' \leq D_{i-1} \leq D_i$	$1 - \frac{M'}{D_{i-1}}$	$\frac{M'}{D_{i-1}} - \frac{M'}{D_i}$	$\frac{M'}{D_i}$
Case 1c $M \leq a_{i+1} \leq J_i$ $M' \leq D_{i-1} = D_i$	$1 - \frac{M'}{D_{i-1}}$	0	$\frac{M'}{D_{i-1}}$
Case 2a $J_i < a_{i+1} \leq M$ $D_{i-1} < D_i \leq M'$	0	0	1
Case 2b $a_{i+1} \leq M \leq J_i$ $M' \leq D_{i-1} = D_i$	$1 - \frac{M'}{D_{i-1}}$	0	$\frac{M'}{D_{i-1}}$
Case 2c $a_{i+1} \leq J_i < M$ $D_{i-1} = D_i \leq M'$	0	0	1

Table 1. Probability of the event $X < b_i$, $X = b_i$ and $X > b_i$ for the six different cases

Case 1b: $[b_i \leq M \leq J_i \leq a_{i+1}]$ As observed earlier, we have that $E[X|X < b_i] \geq \frac{M}{2} \geq \frac{b_i}{2}$ and $E[X|X > b_i] = \frac{M+b_i}{2} \geq b_i$. Setting $x = \frac{M'}{D_{i-1}}$ and using (1) again,

$$\frac{ALG}{OPT} \geq (1-x)\frac{1}{2} + x\left(1 - \frac{D_{i-1}}{D_i} + \frac{D_{i-1}}{D_i}\frac{b_i}{a_{i+1}}\right) =: \beta.$$

To prove $\beta \geq 1/2$ it is enough to prove that $\frac{D_{i-1}}{D_i}\left(1 - \frac{b_i}{a_{i+1}}\right) \leq \frac{1}{2}$. This follows from the fact that $a_{i+1} = b_i + D_i \geq 2D_{i-1}$.

Case 1c: $[b_i \leq M \leq a_{i+1} \leq J_i]$ As before we have that $E[X|X < b_i] \geq \frac{M}{2} \geq \frac{b_i}{2}$ and $E[X|X > b_i] = \frac{M+b_i}{2} \geq \frac{a_{i+1}}{2}$. The last inequality follows because $a_{i+1} \leq J_i = b_i + D_{i-1} \leq b_i + M$. Plugging these back in (1) gives $\frac{ALG}{OPT} \geq \frac{1}{2}$.

Case 2a: $[J_i \leq a_{i+1} \leq M]$ From (2) and Lemma 2, $E[X] = M - \frac{D_i}{2} \geq M/2$ since $M \geq D_i$.

Case 2b: $[a_{i+1} \leq M \leq J_i]$ In this case, it is enough to show that both $E[X|X < b_i]$ and $E[X|X > b_i]$ are bigger than $M/2$. From Lemma 2, $E[X|X < b_i] \geq \frac{M}{2}$. $E[X|X > b_i] = M - \frac{M'}{2} \geq \frac{M}{2}$.

Case 2c: $[a_{i+1} \leq J_i \leq M]$ The analysis is identical to Case 2a.

In fact, the competitive ratio of our algorithm is $1-\epsilon$ if $\epsilon \geq \max\{\frac{D_i-1}{b_i}, \frac{D_i}{a_{i+1}}\}$. The proof is essentially the same as above.

B Proof of Theorem 4

We will now give the detailed proof of Theorem 4.

Let us first fix some notations. Let the optimal single price allocation algorithm for the set S allocates $x^*(S,j)$ items at price $p^*(S,j)$ after j items are produced. So the optimal algorithm generates revenue $OPT(S,i) := x^*(p,i)p^*(S,i)$. Let a fictitious run of the online multi-unit allocation algorithm (that we use as a sub-routine) on the set S allocates $x(S,j)$ items at price $p(S,j)$ after j items are produced and hence generates revenue $ALG(S,i) := x(p,i)p(S,i)$. Recall $\mathbf{B}$ is the set of all the bidders and if M is the number of items finally produced then the optimal revenue against which we compare our algorithm is $OPT(\mathbf{B}, M) =: OPT$. Also to make the notations less messy we will assume M is even (otherwise we will have to carry the floors and ceilings all over the proof and they add little to the understanding of the proof).

Since the online algorithm we use as a subroutine has a competitive ratio of α so for all j,

$$ALG(S,j) \geq \alpha OPT(S,j).$$

Definition 8. *For any price p let $n(S,p)$, $n(T,p)$ and $n(\mathbf{B},p)$ be the number of bids more that p that are made by bidders in S, T and $\mathbf{B}$ respectively.*

Let Y_i be the indicator variable indicating whether the bidder i is in S or not. Let $f_p(Y_1, \ldots, Y_n)$ calculates the number of bids more than or equal to p that are in S, that is $f_p(Y_1, \ldots, Y_n) = n(S,p)$. Note that since the bidders are randomly placed in S or T we have

$$E[f_p(Y_1, Y_2, \ldots, Y_n)] = \frac{n(\mathbf{B},p)}{2}$$

Let c_i is the maximum change in the value of f_p if we change the value of Y_i. Note that c_i is equal to the number of bids of bidder i that are more than p, that is $c_i = n(i,p)$. But from our assumption on the bidder dominance parameter we have $n(i,p) < \eta OPT/p$. Hence

$$\sum c_i^2 < \frac{\eta OPT}{p}\left(\sum c_i\right) = \frac{\eta OPT}{p} n(\mathbf{B},p)$$

By McDiarmid's Inequality we have

$$\Pr\left[\left|\frac{n(\mathbf{B},p)}{2} - n(S,p)\right| > \gamma n(\mathbf{B},p)\right] < \exp\left(\frac{-2\gamma^2 n(\mathbf{B},p)^2}{\sum c_i^2}\right) < \exp\left(\frac{-2p\gamma^2 n(\mathbf{B},p)}{\eta OPT}\right) \tag{3}$$

Lemma 3. *With probability at least* $(1-2|Q|\exp(-2\gamma^2/\eta))$ *both* $OPT(S, M/2)$ *and* $OPT(T, M/2)$ *are greater than* $(1/2 - \gamma)OPT$.

Proof. Let p be any price satisfying $n(\mathbf{B}, M)p \geq OPT$. Then from Equation 3 we have

$$\Pr\left[\left|\frac{n(\mathbf{B},p)}{2} - n(S,p)\right| > \gamma n(\mathbf{B},p)\right] < \exp\left(\frac{-2\gamma^2}{\eta}\right)$$

Now if the optimal algorithm decides to allocate $x^*(\mathbf{B}, M)$ items at price $p^*(\mathbf{B}, M)$ then clearly $n^*(\mathbf{B}, p^*(\mathbf{B}, M))p^*(\mathbf{B}, M) \geq OPT$. Also $p^*(\mathbf{B}, M)$ can take values only from the set Q. So by union bound we have for any M with probability at least $(1 - |Q|exp(-2\gamma^2/\eta))$ we have

$$\left|\frac{n(\mathbf{B}, p^*(\mathbf{B}, M))}{2} - n(S, p^*(\mathbf{B}, M))\right| > \gamma n(\mathbf{B}, p^*(\mathbf{B}, M))$$

That is, with probability at least $(1 - |Q|\exp(-2\gamma^2/\eta))$ at least $(1/2 - \gamma)n(\mathbf{B}, p^*(\mathbf{B}, M))$ bids in S are more than $p^*(\mathbf{B}, M)$. So

$$OPT(S, M/2) \geq n(S, p^*(\mathbf{B}, M))p^*(\mathbf{B}, M) \geq \left(\frac{1}{2} - \gamma\right) OPT.$$

Similarly with probability at least $(1 - |Q|\exp(-2\gamma^2/\eta))$ we have

$$OPT(T, M/2) > \left(\frac{1}{2} - \gamma\right) OPT.$$

Corollary 2. *With probability at least* $(1 - 2|Q|\exp(-2\gamma^2/\eta))$ *we have the following two inequalities*

$$x(S, M/2) > \alpha\left(\frac{1}{2} - \gamma\right)\frac{OPT}{p(S, M/2)}$$

$$x(T, M/2) > \alpha\left(\frac{1}{2} - \gamma\right)\frac{OPT}{p(S, M/2)}.$$

Proof. Since the allocation algorithm that we use as a subroutine has competitive ratio α so,

$$x(S, M/2)p(S, M/2) = ALG(S, M/2) \geq \alpha OPT(S, M/2) > \alpha\left(\frac{1}{2} - \gamma\right)OPT,$$

the last inequality following from Lemma 3.

The other inequality follows similarly.

Corollary 3. *With probability* $(1 - 4|Q|\exp(-2\gamma^2\alpha(1/2 - \gamma)/\eta))$ *for all* $p = p(S, M/2)$ *we have the following two inequalities:*

$$\left|\frac{n(\mathbf{B}, p)}{2} - n(S, p)\right| > \gamma n(S, p)$$

$$\left|\frac{n(\mathbf{B}, p)}{2} - n(T, p)\right| > \gamma n(S, p)$$

and for all $p = p(T, M/2)$ *we have the following two inequalities:*

$$\left|\frac{n(\mathbf{B}, p)}{2} - n(S, p)\right| > \gamma n(T, p)$$

$$\left|\frac{n(\mathbf{B}, p)}{2} - n(T, p)\right| > \gamma n(T, p)$$

Proof. From Equation 3 and Lemma 3 we see that for any fixed $p = p(S, M/2)$

$$\Pr\left[\left|\frac{n(\mathbf{B}, p)}{2} - n(S, p)\right| > \gamma n(\mathbf{B}, p)\right] < \exp\left(\frac{-2\gamma^2 p(n(S, p))}{\eta OPT}\right) < \exp\left(\frac{-2\gamma^2(1/2 - \gamma)}{\eta}\right),$$

the last inequality follows from Corollary 2. The Corollary now follows by applying union bound.

Corollary 4. *With probability* $(1 - 4|Q| \exp(-2\gamma^2\alpha(1/2 - \gamma)/\eta))$

$$n(i, p(S, M/2)) < \gamma x(S, M/2).$$

Proof.

$$n(i, p(S, M/2)) < \eta \frac{OPT}{p(S, M/2)} < \eta \frac{x(S, M/2)}{(1/2 - \gamma)\alpha} < \gamma x(S, M/2),$$

the second inequality follows from Corollary 2 and the third inequality follows from the choice of η.

Now the fictitious run of the online allocation algorithm on S finally decides to allocate $x(S, M/2)$ copies to bidders in S at price $p(S, M/2)$. So the mechanism allocates $(1-6\gamma)x(S, M/2)$ copies to bidders in T. By Corollary 3 with probability at least $(1 - 4|Q| \exp(-2\gamma^2(1/2 - \gamma)/\eta))$

$$n(T, p(S, M/2)) > (1 - 6\gamma)n(S, p(S, M/2)) + 2\gamma n(S, p(S, M/2))$$

Thus there are at least $2\gamma n(S, p(S, M/2))$ losing bid in T that bids at least $p(S, M/2)$. But from Corollary 4 we see that no bidder has more than $\gamma n(S, p(S, M/2))$ bids above $p(S, M/2)$. So by the VCG auction pricing system each winner in T pays at least $p(S, M/2)$ per copy. So the revenue we get from T is at least

$$p(S, M/2)(1 - 6\gamma)x(S, M/2) = ALG(S, M/2)(1 - 6\gamma)$$

Similarly the revenue we get from S is at least

$$ALG(T, M/2)(1 - 6\gamma)$$

So with probability $(1-4|Q| \exp(-2\gamma^2\alpha(1/2-\gamma)/\eta))$ our revenue earned is at least

$$ALG(S, M/2)(1-6\gamma)+ALG(T, M/2)(1-6\gamma) > \alpha(1-6\gamma)(1-2\gamma)OPT > \alpha(1-8\gamma)OPT$$

So if $\gamma = \epsilon/8$ and $(2\gamma^2\alpha(1/2-\gamma)/\eta)) > \log(4|Q|/\delta)$ then with probability $(1 - \delta)$ the total revenue earned on expectation is at least $\alpha(1 - \epsilon)OPT$.