Randomized Online Algorithms for the Buyback Problem

B.V. Ashwinkumar and Robert Kleinberg*

Cornell University, Ithaca, NY
{ashwin85,rdk}@cs.cornell.edu

Abstract. In the matroid buyback problem, an algorithm observes a sequence of bids and must decide whether to accept each bid at the moment it arrives, subject to a matroid constraint on the set of accepted bids. Decisions to reject bids are irrevocable, whereas decisions to accept bids may be canceled at a cost which is a fixed fraction of the bid value. We present a new randomized algorithm for this problem, and we prove matching upper and lower bounds to establish that the competitive ratio of this algorithm, against an oblivious adversary, is the best possible. We also observe that when the adversary is adaptive, no randomized algorithm can improve the competitive ratio of the optimal deterministic algorithm. Thus, our work completely resolves the question of what competitive ratios can be achieved by randomized algorithms for the matroid buyback problem.

1 Introduction

Imagine a seller allocating a limited inventory (e.g. impressions of a banner ad on a specified website at a specified time in the future) to a sequence of potential buyers who arrive sequentially, submit bids at their arrival time, and expect allocation decisions to be made immediately after submitting their bid. An informed seller who knows the entire bid sequence can achieve much higher profits than an uninformed seller who discovers the bids online, because of the possibility that a very large bid is received after the uninformed seller has already allocated the inventory. A number of recent papers [1,2] have proposed a model that offsets this possibility by allowing the uninformed seller to cancel earlier allocation decisions, subject to a penalty which is a fixed fraction of the canceled bid value. This option of canceling an allocation and paying a penalty is referred to as *buyback*, and we refer to online allocation problems with a buyback option as *buyback problems*.

Buyback problems have both theoretical and practical appeal. In fact, Babaioff et al. [1] report that this model of selling was described to them by the ad marketing group at a major Internet software company. Constantin et al. [2] cite numerous other applications including allocation of TV, radio, and newsprint advertisements; they also observe that advance booking with cancellations is a common practice in the airline industry, where limited inventory is

* This research was supported by NSF CAREER Award CCF-0643934.

S. Leonardi (Ed.): WINE 2009, LNCS 5929, pp. 529–536, 2009.

oversold and then, if necessary, passengers are "bumped" from flights and compensated with a penalty payment, often in the form of credit for future flights.

Different buyback problems are distinguished from each other by the constraints that express which sets of bids can be simultaneously accepted. In the simplest case, the only constraint is a fixed upper bound on the total number of accepted bids. Alternatively, there may be a bipartite graph whose two vertex sets are called *bids* and *slots*, and a set of bids may be simultaneously accepted if and only if each bid in the set can be matched to a different slot using edges of the bipartite graph. Both of these examples are special cases of the *matroid buyback problem*, in which there is a matroid structure on the bids, and a set of bids may be simultaneously accepted if and only if they constitute an independent set in this matroid. Other types of constraints (e.g. knapsack constraints) have also been studied in the context of buyback problems [1], but the matroid buyback problem has received the most study. This is partly because of its desirable theoretical properties — the offline version of the problem is computationally tractable, and the online version admits an online algorithm whose payoff is identical to that of the omniscient seller when the buyback penalty is zero — and partly because of its well-motivated special cases, such as the problem of matching bids to slots described above.

As is customary in the analysis of online algorithms, we evaluate algorithms according to their competitive ratio: the worst-case upper bound on the ratio between the algorithm's (expected) payoff and that of an informed seller who knows the entire bid sequence and always allocates to an optimal feasible subset without paying any penalties. The problem of deterministic matroid buyback algorithms has been completely solved: a simple algorithm was proposed and analyzed by Constantin et al. [2,3] and, independently, Babaioff et al. [4], and it was recently shown [1] that the competitive ratio of this algorithm is optimal for deterministic matroid buyback algorithms, even for the case of rank-one matroids (i.e., selling a single indivisible good). However, this competitive ratio can be strictly improved by using a randomized algorithm against an oblivious adversary. Babaioff et al. [1] showed that this result holds when the buyback penalty factor is sufficiently small, and they left open the question of determining the optimal competitive ratio of randomized algorithms — or even whether randomized algorithms can improve on the competitive ratio of the optimal deterministic algorithm when the buyback factor is large.

Our work resolves this open question by supplying a randomized algorithm whose competitive ratio (against an oblivious adversary) is optimal for all values of the buyback penalty factor. We present the algorithm and the upper bound on its competitive ratio in Section 3 and the matching lower bound in Section 4. Our algorithm is also much simpler than the randomized algorithm of [1], avoiding the use of stationary renewal processes. It may be viewed as an online randomized reduction that transforms an arbitrary instance of the matroid buyback problem into a specially structured instance on which deterministic algorithms are guaranteed to perform well. Our matching lower bound relies on defining and analyzing a suitable continuous-time analogue of the single-item buyback problem.

Adaptive adversaries. In this paper we analyze randomized algorithms with an oblivious adversary. If the adversary is adaptive[1], then no randomized algorithm can achieve a better competitive ratio than that achieved by the optimal deterministic algorithm. This fact is a direct consequence of a more general theorem asserting the same equivalence for the class of *request answer games* (Theorem 2.1 of [5] or Theorem 7.3 of [6]), a class of online problems that includes the buyback problem.[2]

Strategic considerations. In keeping with [1,4], we treat the buyback problem as a pure online optimization with non-strategic bidders. For an examination of strategic aspects of the buyback problem, we refer the reader to [2].

Related work. We have already discussed the work of Babaioff et al. [1,4] and of Constantin et al. [2,3] on buyback problems. Prior to this aforementioned work, several earlier papers considered models in which allocations, or other commitments, could be cancelled at a cost. Biyalogorsky et al. [7] studied such "opportunistic cancellations" in the setting of a seller allocating N units of a good in a two-period model, demonstrating that opportunistic cancellations could improve allocative efficiency as well as the seller's revenue. Sandholm and Lesser [8] analyzed a more general model of "leveled commitment contracts" and proved that leveled commitment never decreases the expected payoff to either contract party. However, to the best of our knowledge, the buyback problem studied in this paper and its direct precursors [1,2,3,4] is the first to analyze commitments with cancellation costs in the framework of worst-case competitive analysis rather than average-case Bayesian analysis.

2 Preliminaries

Consider a matroid[3] $(\mathcal{U}, \mathcal{I})$ where $\mathcal{U}$ is the ground set and $\mathcal{I}$ is the set of independent subsets of $\mathcal{U}$. We will assume that the ground set $\mathcal{U}$ is identified with the set $\{1, \ldots, n\}$. There is a bid value $v_i \geq 0$ associated to each element $i \in \mathcal{U}$. The information available to the algorithm at time k $(1 \leq k \leq n)$ consists of the first k elements of the bid sequence — i.e. the subsequence $v_1, v_2, \ldots, v_k$ — and the restriction of the matroid structure to the first k elements. (In other words, for every subset $S \subseteq \{1, 2, \ldots, k\}$, the algorithm knows at time k whether $S \in \mathcal{I}$.)

[1] A distinction between *adaptive offline* and *adaptive online* adversaries is made in [5,6]. When we refer to an adaptive adversary in this paper, we mean an adaptive offline adversary.

[2] The definition of request answer games in [6] requires that the game must have a minimization objective, whereas ours has a maximization objective. However, the proof of Theorem 7.3 in [6] goes through, with only trivial modifications, for request answer games with a maximization objective.

[3] See [9] for the definition of a matroid.

At any step the algorithm can choose a subset $S^k \subseteq S^{k-1} \cup \{k\}$. This set S^k must be an independent set, i.e $S^k \in \mathcal{I}$. Hence the final set held by the algorithm is $R = S^n$. The algorithm must perform a buyback for every element of $B = \left(\cup_{i=1}^n S^i\right) \setminus S^n$. For any set $S \subseteq \mathcal{U}$ let $\text{val}(S) = \sum_{i \in S} v_i$. Finally we define the payoff of the algorithm as $\text{val}(R) - f \cdot \text{val}(B)$.

3 Randomized Algorithm against Oblivious Adversary

This section gives a randomized algorithm with competitive ratio $-W\left(\frac{-1}{e(1+f)}\right)$ against an oblivious adversary. Here W is Lambert's W function[4], defined as the inverse of the function $z \mapsto ze^z$. The design of our randomized algorithm is based on two insights:

1. Although the standard greedy online algorithm for picking a maximum-weight basis of a matroid can perform arbitrarily poorly on a worst-case instance of the buyback problem, it performs well when the ratios between values of different matroid elements are powers of some scalar $r > 1 + f$. (We call such instances "r-structured.")
2. There is a randomized reduction from arbitrary instances of the buyback problem to instances that are r-structured.

3.1 The Greedy Algorithm and r-Structured Instances

Definition 1. *Let $r > 1$ be a constant. An instance of the matroid buyback problem is* r-structured *if for every pair of elements i, j, the ratio v_i/v_j is equal to r^l for some $l \in \mathbb{Z}$.*

Lemma 1. *Consider the greedy matroid algorithm* GMA *that always sells elements and buys them back as necessary to maintain the invariant that the set S^k is a maximum-weight basis of $\{1, 2, \ldots, k\}$. For $r > 1 + f$, when the greedy algorithm is executed on an r-structured instance of the matroid buyback problem, its competitive ratio is at most $\frac{r-1}{r-1-f}$.*

Proof. As is well known, at termination the set S selected by GMA is a maximum-weight basis of the matroid. To give an upper bound on the total buyback payment, we define a set $B(i)$ for each $i \in \mathcal{U}$ recursively as follows: if GMA never sold to i, or sold to i without simultaneously buyback back any element, then $B(i) = \emptyset$. If GMA sold to i while buying back j, then $B(i) = \{j\} \cup B(j)$. By induction on the cardinality of $B(i)$, we find that the set $\{v_x/v_i \mid x \in B(i)\}$ consists of distinct negative powers of r, so

$$\sum_{x \in B(i)} v_x \leq v_i \cdot \sum_{i=1}^{\infty} r^{-i} = \frac{v_i}{r-1}.$$

[4] Lambert's W function is multivalued for our domain. We restrict to the case where $W\left(\frac{-1}{e(1+f)}\right) \leq -1$.

Algorithm Filter(ALG):

Initialize $S = \emptyset$.
for $i = 1, 2, \ldots, n$ **do**
  Observe v_i, w_i.
  Let $x_i = 1$ with probability w_i/v_i, else $x_i = 0$.
  Present i with value w_i to ALG.
  if ALG sells to i **and** $x_i = 1$ **then**
    Sell to i.
  end if
  if ALG buys back j **and** $x_j = 1$ **then**
    Buy back j.
  end if
end for

Algorithm RandAlg(r):

Given: a parameter $r > 1 + f$.
Sample uniformly random $u \in [0, 1]$.
for all elements i **do**
  Let $z_i = u + \lfloor \ln_r(v_i) - u \rfloor$.
  Let $w_i = r^{z_i}$.
end for
Run Filter(GMA) on instances $\mathbf{v}, \mathbf{w}$.

Fig. 1. Randomized algorithms Filter(ALG) and RandAlg(r)

By induction on the number of iterations of the main loop, the set $\bigcup_{i \in S} B(i)$ consists of all the elements ever bought back by GMA; consequently, the total buyback payment is bounded by $f \cdot \sum_{i \in S} \sum_{x \in B(i)} v_x \leq \frac{f}{r-1} \sum_{i \in S} v_i$. Thus, the algorithm's net payoff is at least $1 - \frac{f}{r-1}$ times the value of the maximum weight basis. □

3.2 The Random Filtering Reduction

Consider two instances $\mathbf{v}, \mathbf{w}$ of the matroid buyback problem, consisting of the same matroid $(\mathcal{U}, \mathcal{I})$, with its elements presented in the same order, but with different values: element i has values v_i, w_i in instances $\mathbf{v}, \mathbf{w}$, respectively. Assume furthermore that $v_i \geq w_i$ for all i, and that both values v_i, w_i are revealed to the algorithm at the time element i arrives. Given a (deterministic or randomized) algorithm ALG which achieves expected payoff P on instance $\mathbf{w}$, we present in Figure 1 an algorithm Filter(ALG) achieving expected payoff P on instance $\mathbf{v}$.

Lemma 2. *The expected payoff of* Filter(ALG) *on instance* $\mathbf{v}$ *equals the expected payoff of* ALG *on instance* $\mathbf{w}$*.*

Proof. For each element $i \in \mathcal{U}$, let $\sigma_i = 1$ if ALG sells to i, and let $\beta_i = 1$ if ALG buys back i. Similarly, let $\sigma'_i = 1$ if Filter(ALG) sells to i, and let $\beta'_i = 1$ if Filter(ALG) buys back i. Observe that $\sigma'_i = \sigma_i x_i$ and $\beta'_i = \beta_i x_i$ for all $i \in \mathcal{U}$, and that the random variable x_i is independent of (σ_i, β_i). Thus,

$$\begin{aligned} \mathbf{E}\left[\sum_{i \in \mathcal{U}} \sigma'_i v_i - (1+f)\beta'_i v_i\right] &= \mathbf{E}\left[\sum_{i \in \mathcal{U}} \sigma_i x_i v_i - (1+f)\beta_i x_i v_i\right] \\ &= \sum_{i \in \mathcal{U}} \mathbf{E}[\sigma_i - (1+f)\beta_i]\mathbf{E}[x_i v_i] = \sum_{i \in \mathcal{U}} \mathbf{E}[\sigma_i - (1+f)\beta_i] w_i \\ &= \mathbf{E}\left[\sum_{i \in \mathcal{U}} \sigma_i w_i - (1+f)\beta_i w_i\right]. \end{aligned}$$

The left side is the expected payoff of $\mathsf{Filter}(\mathsf{ALG})$ on instance $\mathbf{v}$ while the right side is the expected payoff of ALG on instance $\mathbf{w}$. □

3.3 A Randomized Algorithm with Optimal Competitive Ratio

In this section we put the pieces together, to obtain a randomized algorithm with competitive ratio $-W\left(\frac{-1}{e(1+f)}\right)$ against oblivious adversary[5]. The algorithm $\mathsf{RandAlg}(r)$ is presented in Figure 1.

Lemma 3. *For all $i \in \mathcal{U}$, we have $v_i \geq w_i$ and $\mathbf{E}[w_i] = \frac{r-1}{r\ln(r)} v_i$.*

Proof. The random variable $\ln_r(v_i) - z_i$ is equal to the fractional part of the number $\ln_r(v_i) - u$, which is uniformly distributed in $[0,1]$ since u is uniformly distributed in $[0,1]$. It follows that w_i/v_i has the same distribution as r^{-u}, which proves that $v_i \geq w_i$ and also that

$$\mathbf{E}\left[\frac{w_i}{v_i}\right] = \int_0^1 r^{-u}\, du = -\frac{1}{\ln(r)} \cdot r^{-u}\Big|_0^1 = \frac{r-1}{r\ln(r)}.$$

□

Theorem 1. *The competitive ratio of* $\mathsf{RandAlg}(r)$ *is* $\frac{r\ln(r)}{r-1-f}$.

Proof. Let $S^* \subseteq \mathcal{U}$ denote the maximum-weight basis of $(\mathcal{U}, \mathcal{I})$ with respect to the weights $\mathbf{v}$. Since the mapping from v_i to w_i is monotonic (i.e., $v_i \geq v_j$ implies $w_i \geq w_j$), we know that S^* is also a maximum-weight basis of $(\mathcal{U}, \mathcal{I})$ with respect to the weights $\mathbf{w}$[6]. Let $v(S^*) = \sum_{i \in S^*} v_i$ and let $w(S^*) = \sum_{i \in S^*} w_i$.

The input instance $\mathbf{w}$ is r-structured, so the payoff of GMA on instance $\mathbf{w}$ is at least $\frac{r-1-f}{r-1} w(S^*)$. The modified weights w_i satisfy two properties that allow application of algorithm $\mathsf{Filter}(\mathsf{ALG})$: the value of w_i can be computed online when v_i is revealed at the arrival time of element i, and it satisfies $w_i \leq v_i$. By Lemma 2, the expected payoff of $\mathsf{Filter}(\mathsf{GMA})$ on instance $\mathbf{v}$, conditional on the values $\{w_i : i \in \mathcal{U}\}$, is at least $\left(\frac{r-1-f}{r-1}\right) \cdot w(S^*)$. Finally, by Lemma 3 and linearity of expectation, $\mathbf{E}\left[w(S^*)\right] \geq \left(\frac{r-1}{r\ln(r)}\right) \cdot v(S^*)$. The theorem follows by combining these bounds. □

The function $f(r) = \frac{r\ln(r)}{r-1-f}$ on the interval $r \in (1+f, \infty)$ is minimized when $-\frac{r}{1+f} = W\left(\frac{-1}{e(1+f)}\right)$ and $f(r) = -W\left(\frac{-1}{e(1+f)}\right)$. This completes our analysis of the randomized algorithm $\mathsf{RandAlg}(r)$.

4 Lower Bound

We prove the lower bound against an oblivious adversary. The proof first reduces to a continuous version of the problem and then applies Yao's Principle [10]. A detailed version of the proof sketches can be found at [11].

---

[5] Note that the algorithm is written in an offline manner just for convenience and can be implemented as an online algorithm.

[6] There may be other maximum-weight basis of $\mathbf{w}$ which were not maximum-weight basis of $\mathbf{v}$.

4.1 Reduction to Continuous Version

Consider a new problem. Time starts at $t = 1$ and stops at time $t = x$, where x is not known to the algorithm. The algorithm at any instant in time can make a mark. The payoff of the algorithm is equal to the time at which it made its final mark minus f times the sum of times of marks before the final mark. We note that any algorithm for the single item buyback problem with competitive ratio c can be transformed into an algorithm for the continuous case with competitive ratio $c \times (1 + \epsilon)$ for arbitrarily small $\epsilon > 0$, by discretizing time into small intervals. We prove lower bound for this new problem.

4.2 Lower Bound against Oblivious Adversaries

Theorem 2. *Any randomized algorithm for the continuous version of the single item buyback problem has competitive ratio at least* $-W\left(\frac{-1}{e(1+f)}\right)$.

The proof is an application of Yao's Principle [10]. We give a one-parameter family of input distributions (parametrized by a number $y > 1$) for the continuous version and prove that any deterministic algorithm for the continuous version of the problem must have a competitive ratio which tends to $-W\left(\frac{-1}{e(1+f)}\right)$ as $y \to \infty$. Note that for the continuous version of the problem input is just stopping time x. For a given $y > 1$, let the probability density for the stopping times be defined as follows.

$$\begin{aligned} \mathbf{f}(x) &= 1/x^2 \text{ if } x < y \\ \mathbf{f}(x) &= \quad 0 \quad \text{if } x > y \end{aligned} \tag{1}$$

Note that the above definition is not a valid probability density function, so we place a point mass at $x = y$ of probability $\frac{1}{y}$. Hence our distribution is a mixture of discrete and continuous probability. For notational convenience let $d(F(x)) = \mathbf{f}(x)$ where F is the cumulative distribution function. Also let $G(x) = 1 - F(x)$. Any deterministic algorithm is defined by a set $T = \{u_1, u_2, \ldots, u_k\}$ of times at which it makes a mark (given that it does not stop before that time).

Lemma 4. *There exists an optimal deterministic algorithm described by the set* $T = \{1, w, w^2, \ldots, w^{k-1}\}$ *for some* w,k.

Proof. Let $T = \{u_1, u_2, \ldots, u_k\}$. We prove that $u_i = u_{i+1}^{(i-1)/i}$ for $i \in [k-1]$ by induction and it is easy to see that the claim follows from this. For lack of space we just prove the inductive case. Let $u_0 = 0$ and $u_{k+1} = \infty$. Let P be the expected payoff of the algorithm.

Note that $P = \sum_{i=1}^{k} \int_{u_i}^{u_{i+1}} (u_i - f \cdot \sum_{j=1}^{i-1} u_j)\, d(F(y))$. We can rewrite the equation as $P = \sum_{i=1}^{k} (u_i - (1+f) \cdot u_{i-1}) \cdot G(u_i)$. If we differentiate P with respect to u_i, equate to 0, and solve, then we obtain the equation $u_i^2 = u_{i-1} \cdot u_{i+1}$. By induction we know that $u_{i-1} = u_i^{(i-2)/(i-1)}$. Substituting and solving we get the necessary equation. □

Lemma 5. *For any algorithm described by* $T = \{1, w, w^2, \ldots, w^{k-1}\}$*, the competitive ratio is bounded below by a number which tends to* $-W\left(\frac{-1}{e(1+f)}\right)$ *as* y *tends to* ∞.

Proof. For lack of space we just give a sketch here. Note that if V is the expected payoff of a prophet who knows the stopping time x, then $V = 1 + \ln(y)$. Also for any algorithm described by $T = \{1, w, w^2, \ldots, w^{k-1}\}$ we have that $P = 1 + (k-1) \cdot \frac{w-1-f}{w}$. Hence if c is the competitive ratio then $c = V/P$. By simple manipulation we see that this is larger than a number which tends to $-W\left(\frac{-1}{e(1+f)}\right)$ as y tends to ∞. □

References

1. Babaioff, M., Hartline, J.D., Kleinberg, R.: Selling ad campaigns: online algorithms with buyback. In: Proc. 10th ACM Conference on Electronic Commerce (2009)
2. Constantin, F., Feldman, J., Muthukrishnan, S., Pál, M.: An online mechanism for ad slot reservations with cancellations. In: Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1265–1274 (2009)
3. Constantin, F., Feldman, J., Muthukrishnan, S., Pál, M.: An online mechanism for ad slot reservations with cancellations. In: Proc. 4th Workshop on Ad Auctions (2008)
4. Babaioff, M., Hartline, J.D., Kleinberg, R.: Online algorithms with buyback. In: Proc. 4th Workshop on Ad Auctions (2008)
5. Ben-David, S., Borodin, A., Karp, R., Tardos, G., Wigderson, A.: On the power of randomization in on-line algorithms. Algorithmica 11(1), 2–14 (1994)
6. Borodin, A., El-Yaniv, R.: Online Computation and Competitive Analysis. Cambridge University Press, Cambridge (1998)
7. Biyalogorsky, E., Carmon, Z., Fruchter, G.E., Gerstner, E.: Research note: Overselling with opportunistic cancellations. Marketing Science 18(4), 605–610 (1999)
8. Sandholm, T., Lesser, V.R.: Leveled commitment contracts and strategic breach. Games and Economic Behavior 35, 212–270 (2001)
9. Oxley, J.: Matroid Theory. Oxford University Press, Oxford (1992)
10. Yao, A.C.C.: Probabilistic computations: Toward a unified measure of complexity. In: Proceedings of the 18th Annual Symposium on Foundations of Computer Science, Washington, DC, USA, pp. 222–227. IEEE Computer Society, Los Alamitos (1977)
11. Ashwinkumar, B.V., Kleinberg, R.: Randomized online algorithms for the buyback problem. CoRR abs/0908.0043 (2009)