

Online Budgeted Maximum Coverage

Dror Rawitz^{*1} and Adi Rosén^{†2}

1 Faculty of Engineering, Bar-Ilan University, Ramat Gan, Israel
dror.rawitz@biu.ac.il

2 CNRS and Université Paris Diderot, Paris, France
adiro@liafa.univ-paris-diderot.fr

Abstract

We study the ONLINE BUDGETED MAXIMUM COVERAGE (OBMC) problem. Subsets of a weighted ground set U arrive one by one, where each set has a cost. The online algorithm has to select a collection of sets, under the constraint that their cost is at most a given *budget*. Upon arrival of a set the algorithm must decide whether to accept or to reject the arriving set, and it may also drop previously accepted sets (preemption). Rejecting or dropping a set is irrevocable. The goal is to maximize the total weight of the elements covered by the sets in the chosen collection.

We present a deterministic $\frac{4}{1-r}$ -competitive algorithm for OBMC, where r is the maximum ratio between the cost of a set and the total budget. Building on that algorithm, we then present a randomized $O(1)$ -competitive algorithm for OBMC. On the other hand, we show that the competitive ratio of any deterministic online algorithm is $\Omega(\frac{1}{\sqrt{1-r}})$.

We also give a deterministic $O(\Delta)$ -competitive algorithm, where Δ is the maximum weight of a set (given that the minimum element weight is 1), and if the total weight of all elements, $w(U)$, is known in advance, we show that a slight modification of that algorithm is $O(\min\{\Delta, \sqrt{w(U)}\})$ -competitive. A matching lower bound of $\Omega(\min\{\Delta, \sqrt{w(U)}\})$ is also given.

Previous to the present work, only the unit cost version of OBMC was studied under the online setting, giving a 4-competitive algorithm [36]. Finally, our results, including the lower bounds, apply to REMOVABLE ONLINE KNAPSACK which is the preemptive version of the ONLINE KNAPSACK problem.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases budgeted coverage, maximum coverage, online algorithms, competitive analysis, removable online knapsack

Digital Object Identifier 10.4230/LIPIcs.ESA.2016.73

1 Introduction

The BUDGETED MAXIMUM COVERAGE problem (abbreviated BMC) is the dual of the classical SET COVER problem. In SET COVER the input consists of a collection of sets $\mathcal{S} = \{S_1, \dots, S_m\}$ over the ground set $U = \{u_1, \dots, u_n\}$ with cost for each set, and the goal is to find a sub-collection of sets of minimal cost, whose union covers all the elements in the ground set. In BMC we are, in addition, given weights for the elements and a budget cap B . The goal is to find a subcollection of the sets that maximizes the total weight of the union

* Supported in part by a grant from the Israeli Ministry of Science, Technology, and Space, Israel (French-Israeli project Maimonide 31768XL) and by the Israel Science Foundation (grant No. 497/14).

† Research supported in part by ANR project NeTOC, and by a French-Israeli grant PHC Maimonide 31768XL.



of those sets, under the constraint that the total cost of the subcollection is at most B . In other words, the goal is to find a subcollection that maximizes coverage given a knapsack constraint. In fact the KNAPSACK problem can be viewed as the special case of BMC in which the sets are pairwise disjoint.

In the ONLINE BUDGETED MAXIMUM COVERAGE problem (OBMC) sets arrive online and the goal of an online algorithm for this problem is to find (in an online manner) a sub-collection of the sets, that maximizes the weight of the covered items, while adhering to the budget constraint. Preemption of previously used sets is allowed, but a preempted (or rejected at arrival) set cannot be used again later. Preemption is necessary in order to achieve an unbounded competitive ratio, for this problem, by a deterministic or randomized online algorithm. BMC has many applications both in its offline and online versions, such as facility location [32, 8, 7], where the budget is the number of facilities, and web streaming [36], where the budget is the number (or total size) of web objects that can be stored in memory.

The BUDGETED MAXIMUM COVERAGE problem, as well as its dual problem, the SET COVER problem, both in their weighted (general costs) and unweighted (unit costs) versions, are fundamental, widely studied problems (cf. [23, 40]). Even in their unweighted versions they are NP-hard problems [20], while approximation algorithms do exist for their weighted versions (with approximation ratios of $O(\log n)$ for SET COVER [13] and $\frac{e}{e-1}$ for BMC [28]). These approximation ratios are best possible unless $P = NP$ [18, 3]. The online version of the SET COVER problem has been studied in many variants (see, e.g., [15, 2, 11, 29]). As to the OBMC problem, only the unweighted case has been studied in the online setting, where a 4-competitive deterministic algorithm is given [36].

In the present paper we give the first results for the *budgeted* (i.e., when sets have varying costs) online maximum coverage problems. We give both upper and lower bounds on the competitive ratio of deterministic and randomized algorithms for this problem, in terms of a number of parameters of the instance.

1.1 Our Contributions

We present a deterministic $\frac{4}{1-r}$ -competitive algorithm, where $r \triangleq \max_i \frac{c(S_i)}{B}$ is the maximum fraction of the budget needed for any single set. Our algorithm is inspired by the online algorithms for the case of unit costs [36, 4]. However, several new ideas are needed to cope with general costs, such as working with a fractional solution in the background and rounding it in an online manner to an integral solution while incurring only a small penalty. Furthermore, the natural algorithm, that results from reducing the weighted (set costs) case to the unit cost case by duplicating the sets, does not necessarily yield fractional solutions that can be readily converted to integral ones (i.e., many sets could be used only fractionally). Instead, we give an online algorithm for the weighted fractional setting that computes a solution which has at most one set used fractionally. Such solution can be converted to an integral one in an online manner while incurring only a small penalty (a $1/(1-r)$ factor).

On the negative side, we show that the competitive ratio of any deterministic online algorithm for OBMC must depend on r , by showing a lower bound of $\Omega(\frac{1}{\sqrt{1-r}})$. Building on our deterministic algorithms we then also give an $O(1)$ -competitive randomized algorithm for OBMC.

We further give a deterministic $(\Delta + 2)$ -competitive algorithm for OBMC, where $\Delta \triangleq \max_{S \in \mathcal{S}} w(S)$ is the maximum weight of a set (defined under the assumption that all element weights are at least 1). Note that for unit weights we have that Δ is the maximum set size. If $w(U)$, i.e., the total weight of all elements in the ground set, is known in advance, we show that a slight modification of that algorithm is $O(\min\{\Delta, \sqrt{w(U)}\})$ -competitive. We

give a matching lower bound, namely that the competitive ratio of any deterministic online algorithm for OBMC is $\Omega(\min\{\Delta, \sqrt{w(U)}\})$, even for the special case of unit weights. Note that $w(U) = n$ in the unit weights case.

We note that by applying our deterministic upper bounds to the special case of OBMC in which the sets are pairwise disjoint, we obtain results for the REMOVABLE ONLINE KNAPSACK problem. This problem is the version of ONLINE KNAPSACK in which preemption is allowed. Furthermore, our deterministic lower bounds apply to this problem since they can be obtained using constructions containing pairwise disjoint sets.

Due to lack of space some of the proofs are omitted from this extended abstract.

1.2 Related Work

In the MAXIMUM COVERAGE problem the goal is, given an integer parameter k , to cover as many elements as possible, using at most k sets. In this case the natural greedy algorithm computes solutions whose weight is within a factor of $1 - (1 - \frac{1}{k})^k > 1 - \frac{1}{e}$ from the optimum (see [33, 24, 23]). This ratio holds even in the more general case of nonnegative, nondecreasing, submodular set function maximization [34, 19].¹ Khuller, Moss and Naor [28] showed that MAXIMUM COVERAGE cannot be approximated to within a factor better than $\frac{e}{e-1}$, unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$. Feige [18] did the same under the weaker assumption of $\text{P} \neq \text{NP}$. Ageev and Sviridenko [1] presented an approximation algorithm for MAXIMUM COVERAGE that computes solutions whose weight is within a factor of $1 - (1 - \frac{1}{\Delta})^\Delta$ from the optimum, where Δ is the maximum size of a set. Buchbinder et al. [9] studied submodular maximization with cardinality constraints which contain MAXIMUM COVERAGE as a special case. Khuller et al. [28] showed that BMC can be approximated to within $\frac{e}{e-1}$. Sviridenko [38] extended this result to maximization of a monotone submodular set function subject to a budget constraint.

Saha and Getoor [36] presented a deterministic 4-competitive algorithm for the ONLINE MAXIMUM COVERAGE problem. Ausiello et al. [4] analyzed a variant of the above algorithm and showed that its competitive ratio is strictly less than 4, but that it tends to 4 as k increases. They also considered the special case of ONLINE MAXIMUM COVERAGE in which vertices are used to cover edges (i.e., an element appears in exactly two sets) and provided a simple deterministic 2-competitive algorithm for the latter that simply chooses the k largest sets seen so far. Ausiello et al. [4] also gave lower bounds 2 and $\frac{3}{2}$ for ONLINE MAXIMUM COVERAGE and for that special case, respectively, on the competitive ratio of any deterministic online algorithm. Ashwinkumar [39] and Chakrabarti and Kale [12] presented streaming 4-approximation algorithms for maximizing a monotone submodular function subject to cardinality constraint. Buchbinder, Feldman, and Schwartz [10] provided constant competitive ratio algorithms for online submodular maximization with preemption and a cardinality constraint. They also gave a deterministic 4-competitive algorithm for the monotone case.

We note that Awerbuch et al. [5] studied a problem they called ONLINE SET COVER. However they actually consider a variant of ONLINE MAXIMUM COVERAGE in which the elements arrive in an online manner, and the sets are revealed during this process. The goal is to cover as many elements as possible using k sets without preemption, where an element is considered covered only by a set that contains it which is added to the solution after the

¹ A function f is called submodular if $f(T) + f(T') \geq f(T \cup T') + f(T \cap T')$ for every two sets T and T' in the domain of f .

arrival of the element. Awerbuch et al. [5] gave a randomized $O(\log n \log \frac{m}{k})$ -competitive algorithm for this problem.

The dual of MAXIMUM COVERAGE is the classical SET COVER problem. For the unweighted SET COVER problem, Johnson [27] and Lovász [30] showed that the greedy algorithm is an H_n -approximation algorithm, where H_n the n th harmonic number. This result was generalized by Chvátal [13] to the weighted case. Feige [18] proved a lower bound of $(1 - o(1)) \ln n$ on the approximability of that problem (unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$). In [35, 3] it was shown that SET COVER cannot be approximated within a factor of $c \log n$, for some $c > 0$, unless $\text{P} = \text{NP}$. SET COVER can also be approximated to within a factor of $\Delta_U \triangleq \max_{u \in U} |\{S : u \in S\}|$ [22, 6]. However, it is NP-hard to approximate it within $\Delta_U - 1 - \varepsilon$, for any $\varepsilon > 0$, assuming $\Delta_U > 2$ [16], or within 1.36 for $\Delta_U = 2$ [17].

A certain online version of SET COVER was studied by Alon et al. [2]. In this problem the sets are known in advance, subsets of the elements arrive in an online manner, and the goal is to cover all seen elements with a sub collection of sets of minimal cardinality. A deterministic $O(\log m \log n)$ -competitive algorithm and a nearly matching lower bound are given in [2] (n is the number of elements and m the number of sets).

KNAPSACK is a special case of BMC in which the sets are pairwise disjoint. KNAPSACK is known to be NP-hard, but admits an FPTAS [37, 25]. REMOVABLE ONLINE KNAPSACK (ROK) is a special case of OBMC in which the sets are pairwise disjoint. In other words, in ROK items arrive one by one, each with its load and value. An online algorithm is required to accept or to reject an incoming item upon arrival, and it is allowed to drop previously accepted items to make room for a new item. The goal is to maximize the value accrued by the accepted items, under the constraint that their total load is within a given maximum load. Iwama and Taketomi [26] considered the special case of ROK in which the value of an item is equal to its load. They provided a deterministic competitive algorithm whose ratio is $\frac{\sqrt{5}+1}{2} \approx 1.62$, and a matching lower bound. Han, Kawase, and Makino [21] and Cygan, Jež, and Sgall [14] gave a randomized 2-competitive algorithm and showed that the competitive ratio of any randomized online algorithm is at least $\frac{e+1}{e}$. Both lower bounds apply to OBMC. NON-REMOVABLE ONLINE KNAPSACK is the variant in which accepted items cannot be dropped. In this case the deterministic [31] and randomized [41] competitive ratios are known to be unbounded.

1.3 The Model

An instance of the BMC problem is composed of a weighted ground set $U = \{u_1, \dots, u_n\}$, with each element having a known weight $w(u_i) \geq 1$ (we define all weights to be at least 1 to avoid arbitrary scaling of the weights). The instance is further composed of a collection $\mathcal{S} = \{S_1, \dots, S_m\}$ of sets, where $S_i \subseteq U$, for every i . The cost of a set S_i is denoted $c(S_i)$. W.l.o.g. we assume that the budget cap is 1, and therefore $0 < c(S_i) \leq 1$, for every i .

In OBMC the sets of \mathcal{S} arrive online, where each set S_i is given by the elements that it contains, as well as its cost, $c(S_i)$. When a set arrives the online algorithm has to decide whether to *accept* it or to *reject* it, under the constraint that the currently accepted sets at any given time should have a cumulative cost not larger than 1. The algorithm can *drop* a previously accepted set, i.e., extract it from the currently accepted sets. However, a rejected or dropped set cannot later be re-accepted. The goal of the online algorithm is to maximize the total weight of the elements covered by the accepted sets.

Given a set $S \subseteq U$, we define its weight to be $w(S) = \sum_{u \in S} w(u)$. Given a sub-collection $\mathcal{C} \subseteq \mathcal{S}$, we define its cost to be $c(\mathcal{C}) = \sum_{S \in \mathcal{C}} c(S)$. Given an instance of the OBMC problem we define $\Delta \triangleq \max_{S \in \mathcal{S}} w(S)$, i.e., the maximum weight of a set. Note that for unit weights we have $\Delta = \max_{S \in \mathcal{S}} |S|$. Further we define $r \triangleq \max_i c(S_i)$ to be the maximum cost of a set.

2 Deterministic Lower Bound

In this section we give lower bounds on the competitive ratio of deterministic online algorithms for OBMC in terms of three parameters: the number of elements n , the weight of the heaviest set, Δ , and the maximum cost of a set, r .

We start with our lower bound in terms of n and Δ .

► **Theorem 1.** *The competitive ratio of any deterministic online algorithm for OBMC is $\Omega(\min\{\sqrt{n}, \Delta\})$.*

Proof. Let ALG be any deterministic online algorithm for OBMC. Consider an input sequence containing a collection of subsets of a ground set U that contains (at most) k^2 unit weight elements, for an arbitrary integer k . We define the input sequence given by an adversary. The input sequence starts with a set S_0 , where $S_0 = \{u_1, \dots, u_k\}$ and $c(S_0) = 1$. If ALG does not accept S_0 , then the sequence terminates. Otherwise, the sequence continues with S_1, S_2, \dots , such that $S_i = \{u_i\}$ and $c(S_i) = 1/k^2$, for $i \geq 1$, until either ALG drops S_0 or $i = k^2$.

To analyze the competitive ratio of ALG, note that there are three options as to the actual input sequence given by the adversary:

- If ALG rejects S_0 , then OPT covers k elements using S_0 while ALG covers nothing.
- If ALG accepts S_0 and when S_i , for some $i \geq 1$, arrives, ALG drops S_0 and (possibly) accepts S_i , then OPT covers k elements using S_0 , while ALG covers at most one element.
- If ALG accepts S_0 and never drops it, then ALG covers k elements, while OPT covers k^2 elements using S_1, \dots, S_{k^2} .

Hence the competitive ratio of ALG is at least k , and the theorem follows, since $k = \sqrt{n}$ and $k = \Delta$. ◀

A similar construction works for r as well.

► **Theorem 2.** *The competitive ratio of any deterministic online algorithm for OBMC is $\Omega(1/\sqrt{1-r})$. In particular, if $r = 1$ the competitive ratio of any deterministic algorithm is unbounded.*

Proof. The theorem clearly holds for $r \leq \frac{8}{9}$, so we assume for the rest of the proof that $r > \frac{8}{9}$. If $r < 1$, let k be a positive integer such that $\frac{1}{2} \cdot \frac{1}{\sqrt{1-r}} \leq k \leq \frac{1}{\sqrt{1-r}}$. There exists such a k , since $\frac{1}{2} \cdot \frac{1}{\sqrt{1-r}} \geq 1.5$. Also, note that $r + \frac{1}{k^2} \geq r + (1-r) > 1$ and that $\frac{1}{k^2} \leq 4(1-r) \leq \frac{4}{9} < r$. We use the same adversary that is used in Theorem 1 with k defined as above, with $c(S_0) = r$, and with $c(S_i) = 1/k^2$, for $i \geq 1$. The rest of the proof is similar to the proof of Theorem 1, showing that the competitive ratio is $k = \Theta(1/\sqrt{1-r})$.

If $r = 1$, we can pick an arbitrarily large positive integer k , which then shows that the competitive ratio is unbounded. ◀

We can conclude with the following theorem.

► **Theorem 3.** *Let ALG be a c -competitive deterministic online algorithm for OBMC. Then $c = \Omega(\min\{\sqrt{n}, \Delta, \frac{1}{\sqrt{1-r}}\})$.*

We note that the above constructions can be slightly modified to consist of pairwise disjoint sets.

► **Theorem 4.** *Let ALG be a c -competitive deterministic online algorithm for REMOVABLE ONLINE KNAPSACK. Then $c = \Omega(\min\{\sqrt{n}, \Delta, \frac{1}{\sqrt{1-r}}\})$.*

3 $O(\frac{1}{1-r})$ -competitive Algorithm

In this section we present a deterministic $\frac{4}{1-r}$ -competitive algorithm for OBMC. In what follows we assume that $r < 1$. Otherwise, the competitive ratio of any deterministic algorithm is unbounded (see Theorem 2).

Roughly speaking, our algorithm is inspired by the online algorithm for the case of unit costs [36]. We use a similar greedy rule: a set joins the solution if its marginal benefit, with respect to the current solution, is high enough. However, in contrast to the unit cost algorithm, we sort the sets in the current solution by cost effectiveness (to be defined later), and consider only sets that are contained in the prefix of the sets of the current solution which sums to a cost of at most 1. This prefix may be fractional, namely there may be a set that is only partly considered, which then complicates both the algorithm and the analysis. In what follows we define some notations and then present formally the algorithm.

Definitions and notations

Our algorithm is defined based on an imaginary fractional solution to the problem that we maintain throughout receiving the input. In this fractional solution, the algorithm can use only a fraction $x(S) \in [0, 1]$ of a given set S , paying only $x(S) \cdot c(S)$, and covering by set S at most an $x(S)$ fraction of every element $v \in S$. This fractional solution can be defined using the following variables. In what follows we refer by *time* i to the time *after* the algorithm has processed the i th input set. A variable with a subscript i refers to the value of the variable at time i .

- $x_i(S) \in [0, 1]$, for $S \in \mathcal{S}$, is the fraction of set S used by the algorithm at time i .
- $z_i(v, S) \in [0, 1]$, for $S \in \mathcal{S}$ and $v \in U$, is the fraction of v that is covered by the algorithm using set S at time i . We set $z_i(v, S) = 0$ if $v \notin S$.

We further define, given any \vec{z}_i , for $v \in U$, $\hat{z}_i(v) \triangleq \sum_{S \in \mathcal{S}} z_i(v, S)$.

The (fractional) optimization problem can now be defined by the following linear program:

$$\begin{array}{ll}
 \max & \sum_{v \in U} \hat{z}(v) \cdot w(v) \\
 \text{s.t.} & \sum_{S \in \mathcal{S}} x(S) \cdot c(S) \leq 1 \\
 & \hat{z}(v) \leq 1 \quad \forall v \in U \\
 & z(v, S) \leq x(S) \quad \forall v \in U, S \in \mathcal{S}, \\
 & x(S) \in [0, 1] \quad \forall S \in \mathcal{S} \\
 & z(v, S) \geq 0 \quad \forall v \in U, S \ni v \\
 & z(v, S) = 0 \quad \forall v \in U, S \not\ni v
 \end{array}$$

Before presenting the algorithm we need the following further notations.

- $w(\vec{z}) \triangleq \sum_v \hat{z}(v)w(v)$, i.e., the total weight covered in a solution defined by the matrix \vec{z} .
- $\rho(\vec{z}, \vec{x}, S) \triangleq \frac{\sum_v z(v, S)w(v)}{x(S)c(S)}$, if $x(S) > 0$, and $\rho(\vec{z}, \vec{x}, S) \triangleq 0$, otherwise. That is, $\rho(\vec{z}, \vec{x}, S)$ stands for the total weight covered by set S in a solution defined by the matrix \vec{z} , divided by the “actual cost” paid for S . We call this quantity the *efficiency* of set S with respect to the solution (\vec{z}, \vec{x}) .

3.1 The Algorithm

Our algorithm maintains two variables z and x , corresponding to the variables by the same names described above, and which represent a current fractional (imaginary) solution held by the online algorithm. As the algorithm is an online algorithm, we allow it to increase $z(\cdot, S)$ and $x(S)$ only when set S arrives.

Algorithm 1: $\text{insert}(S, x, z)$

```

1  $x' \leftarrow x; z' \leftarrow z$ 
2  $x'(S) \leftarrow 1$ 
3 foreach  $v \in U$  do  $z'(v, S) \leftarrow \begin{cases} 1 - \hat{z}(v) & v \in S, \\ 0 & v \notin S \end{cases}$ 
4  $\hat{S} \leftarrow \{S : x'(S) > 0\}; \ell \leftarrow |\hat{S}|$ 
5 Order the sets in  $\hat{S}$  by non-increasing value of  $\rho(\vec{z}', \vec{x}', S)$ ; let  $S_{j_1}, S_{j_2}, \dots, S_{j_\ell}$  be the
   ordering.
6  $k \leftarrow \max \left\{ k' \leq \ell : \sum_{i=1}^{k'-1} x'(S_{j_i})c(S_{j_i}) < 1 \right\}$ 
7  $\chi \leftarrow \min \left\{ \frac{1 - \sum_{i=1}^{k-1} x'(S_{j_i})c(S_{j_i})}{c(S_{i_k})}, x'(S_{i_k}) \right\}$ 
8 foreach  $v \in S_{i_k}$  do  $z'(v, S_{i_k}) \leftarrow \frac{\chi}{x'(S_{i_k})} \cdot z'(v, S_{i_k})$ 
9  $x'(S_{i_k}) \leftarrow \chi$ 
10 for  $i = k + 1$  to  $\ell$  do
11    $x'(S_{j_i}) \leftarrow 0$ 
12   foreach  $v \in U$  do  $z'(v, S_{j_i}) \leftarrow 0$ 
13 return  $(x', z')$ 

```

Algorithm 2: α -greedy; operations when set S_i arrives.

```

1  $x'(S_i) \leftarrow 1$ 
2  $z'(v, S_i) \leftarrow \begin{cases} 1 - \hat{z}(v) & \text{if } v \in S_i \\ 0 & \text{otherwise} \end{cases}$ 
3 if  $\rho(\vec{z}', \vec{x}', S_i) > \alpha \cdot w(z)$  then  $(x, z) \leftarrow \text{insert}(S_i, x, z)$ 

```

We first define a procedure **insert** that we use in the algorithm. This procedure takes a set S and inserts it into the current solution represented by the variables z and x . This changes the values of z and x to represent the new solution.

We can now define the online algorithm, that we call α -greedy, for any $\alpha > 1$. The optimal value for α will be defined later in the analysis.

 α -greedy

We initialize the two (vector) variables $\vec{z} \leftarrow \vec{0}$, $\vec{x} \leftarrow \vec{0}$. Then, for every set S_i that arrives, we use the operations defined in the pseudocode in Algorithm 2.

At any given time, the solution held by the online (regular, integral) algorithm consists of all the sets S for which $x(S) = 1$. As we later prove, the algorithm has, at any given time, at most one set S "used fractionally", i.e., with $x(S) \in (0, 1)$.

We claim that the algorithm is a well defined online algorithm for our problem. That is, that (1) the algorithm accepts a set only when this set arrives, i.e., a set that is not accepted when it arrives, or accepted but subsequently dropped, cannot later be part of the solution; and (2) the solution held by the algorithm at any given time is feasible, i.e., the total budget used by the algorithm is at most 1 at any given time. To see these two points observe that all the changes in the variables held by the algorithm are done in procedure **insert**. Procedure **insert** assigns a value of 1 to variable $x'(S)$, only for S which is the inserted set: an explicit assignment of 1 is only done in Line 1, and in Line 9 the value of $x'(S_{i_k})$ cannot

grow compared to the previous round. This proves point (1). Point (2) requires a bit more of formalism, which is given in the proof of the following lemma.

► **Lemma 5.** *For any time i , $\sum_{j \in O_i} c(S_j) \leq 1$, where $O_i = \{j : x_i(S_j) = 1\}$.*

Proof. We prove this lemma by induction on i . For the base of the induction, i.e., when $i = 0$, we have that $\vec{x}_0 = \vec{0}$. It follows that $O_0 = \emptyset$ and we are done. For the inductive step, we assume that the claim is true for $i - 1$, for $i \geq 1$ and prove it for i . If **insert** is not called during iteration i , then $x_i = x_{i-1}$ and $O_i = O_{i-1}$, and the claim follows from the inductive hypothesis. If **insert** is activated, then x_i is constructed such that $x_i(S_{j_i}) = 0$ if $j_i > k$ and $\sum_{i=1}^k x_i(S_{j_i})c(S_{j_i}) \leq 1$. Hence $\sum_{j \in O_i} c(S_j) \leq 1$. ◀

We now give two technical claims which we use in the analysis. The first claim is immediate from the code of **insert**.

► **Claim 6.** *The efficiency of a given set S remains the same throughout the period when $x(S) > 0$.*

The next claim says that at any time i , there may be at most one set S such that $x(S) \in (0, 1)$, and that if such set exists then the whole budget is used. Furthermore, if such set exists then that set is the one with minimum efficiency among the sets with non-zero x -coefficient. We note that these properties of our algorithm are the properties that allow one to obtain in an online manner an integral solution, and that a simple, natural reduction of the weighted fractional case to the unweighted case does not yield an algorithm with such properties.

► **Claim 7.** *Let $\mathcal{S}_i^+ = \{S : x_i(S) > 0\}$. There is at most one set $S \in \mathcal{S}^+$ such that $x(S) < 1$. If such a set S exists then (1) $\sum_{S \in \mathcal{S}_i^+} x_i(S)c(S) = 1$; and (2) $\rho(\vec{z}_i, \vec{x}_i, S) \leq \rho(\vec{z}_i, \vec{x}_i, S')$, for any $S' \in \mathcal{S}^+ \setminus \{S\}$.*

Proof. We prove the claim by induction on i . For the basis of the induction, i.e., for $i = 0$, the claim is trivial in the empty sense. We now assume that the claim holds for $i - 1$, for $i \geq 1$, and we prove it for i .

If **insert** is not invoked during iteration i , then the claim clearly holds by the induction hypothesis. If procedure **insert** is invoked we have that (according to Line 3 of α -greedy)

$$\rho(z', x', S_i) > \alpha \cdot w(z_{i-1}) = \alpha \cdot \sum_{S \in \mathcal{S}_{i-1}^+} \rho(\vec{z}_{i-1}, \vec{x}_{i-1}, S) \cdot x_{i-1}(S) \cdot c(S).$$

We now consider two cases. The first case is when a set S with $x_{i-1}(S) \in (0, 1)$ exists. In that case, by the induction hypothesis we have that $\sum_{S \in \mathcal{S}_{i-1}^+} x_{i-1}(S)c(S) = 1$. It follows that $\rho(z', x', S_i) > \min_{S \in \mathcal{S}_{i-1}^+} \{\rho(\vec{z}_{i-1}, \vec{x}_{i-1}, S)\}$ (recall that $\alpha > 1$), and therefore the new set S_i is not the last set in the non-increasing order of efficiency defined in Line 5 of procedure **insert**. Points (1) and (2) for time i therefore follow from the code of **insert** and the induction hypothesis of point (2). The second case is when there is no set S with $x_{i-1}(S) \in (0, 1)$. In that case the code of **insert** directly guarantees both points (1) and (2) for time i . ◀

3.2 Competitive Analysis

We analyze the competitive ratio of the algorithm using a charging scheme argument. We first describe the scheme in general terms – more details are given in the following paragraphs. Let OPT be an optimal solution, i.e., an optimal collection of sets. As the sets of OPT arrive,

each new element u that is covered by OPT is allotted $w(u)$ monetary units (or simply money). We show that the allotted money can be moved between elements of U such that all the money is allotted to the items covered by ALG , and such that each element v that is covered by ALG has at most $\frac{4}{1-r} \cdot w(v)$ amount of money. This gives an upper bound of $\frac{4}{1-r}$ on the competitive ratio.

We now formally define the charging scheme. Let $\text{OPT}_k = \text{OPT} \cap \{S_1, \dots, S_k\}$. If $S_i \in \text{OPT}$, let $F_i \triangleq S_i \setminus \cup_{S \in \text{OPT}_{i-1}} S$, i.e., F_i contains all the elements covered by OPT that, when S_i arrives, are covered for the first time by OPT (according to the order of arrival of the sets). When a set $S_i \in \text{OPT}$ arrives, each element $u \in F_i$ is allotted $w(u)$ money. The money that is allotted this way is always held by items covered by ALG , and may sometimes be moved between items covered by ALG ; furthermore this money is partitioned into *blue money* and *red money*. We will denote by $\text{blue}_i(v)$ and by $\text{red}_i(v)$ the amount of blue and red money, respectively, held by $v \in U$ at time i .

Now, the rules that govern the allotment and movement of the money are the following rules, applied when a set S_i arrives.

- **Transfer of money.** This rule is applied if the then part of Line 3 of α -**greedy** is reached, i.e., if procedure **insert** is invoked. In this case elements that lost coverage relinquish part of their money, and this money is transferred to elements that gain coverage. Let $\hat{z}'_i(v) \triangleq \sum_{j=1}^{i-1} z_i(v, S_j)$ and $Z_i \triangleq \{v : \hat{z}'_i(v) < \hat{z}_{i-1}(v)\}$. That is, we look at the coverage of an item v by all but the last arriving set, and observe if this coverage reduced during the course of iteration i .
 1. **Out-transfer of money.** For each $v \in Z_i$ let $\delta_v = (1 - \frac{\hat{z}'_i(v)}{\hat{z}_{i-1}(v)})$. We remove from v an amount of R_v red money which is a δ_v -fraction of its current red money, and an amount of B_v blue money which is a δ_v -fraction of its current blue money.
 2. **In-transfer of money.** These total amounts of removed red and blue money are distributed to the various red and blue money variables of $u \in S_i$ proportionally to $z_i(u, S_i) \cdot w(u)$. That is, each element $u \in S_i$ gets additional $\frac{z_i(u, S_i) \cdot w(u)}{\sum_{v \in S_i} z_i(v, S_i) \cdot w(v)}$ $\sum_{v \in Z_i} R_v$ red money, and gets additional $\frac{z_i(u, S_i) \cdot w(u)}{\sum_{v \in S_i} z_i(v, S_i) \cdot w(v)}$ $\sum_{v \in Z_i} B_v$ blue money.
- **Creation of new money.** If $S_i \in \text{OPT}$, $w(v)$ money is distributed for each $v \in F_i$ as follows: (1) v gets $\hat{z}_i(v)w(v)$ newly created blue money; and (2) a total amount of $R_i(v) = (1 - \hat{z}_i(v))w(v)$ red money is created, and distributed among the items u currently covered by the algorithm, i.e., each element $u \in U$ gets additional red money in the amount of $R_i(v) \frac{\hat{z}_i(u)w(u)}{w(\hat{z}_i)}$ (i.e., the red money is distributed proportionally to the contribution of each element to the total weight of the current solution).

We now prove upper bounds on the amount of red and blue money held by any element at any given time. We start with a technical claim. In the next lemma we show that if S_i is accepted, then the amount of coverage provided by set S_i is more than α times the coverage lost due to the sets (or part of sets) pushed out.

► **Lemma 8.** *Assume that $x_i(S_i) > 0$. Then,*

$$\rho(\vec{z}_i, \vec{x}_i, S_i) > \alpha \cdot \frac{\sum_{j < i} \sum_{u \in U} w(u) [z_{i-1}(u, S_j) - z_i(u, S_j)]}{x_i(S_i)c(S_i)}.$$

Proof. Since $x_i(S) > 0$ we know that $\rho(\vec{z}_i, \vec{x}_i, S_i) = \frac{\sum_{u \in U} (1 - \hat{z}_{i-1}(u))w(u)}{c(S_i)} > \alpha \cdot w(z_{i-1})$ and that **insert** was invoked. Consider what happens to the solution $(\vec{z}_{i-1}, \vec{x}_{i-1})$ during the invocation of **insert** at iteration i . Some sets do not lose coverage (i.e., sets S for

which $x_i(S) = x_{i-1}(S) = 1$). Other sets may leave the cover and their cover is lost (i.e., $x_{i-1}(S) > x_i(S) = 0$). In addition, by Claim 7, there may be at most one set that partly leaves the cover, and so it both retains and loses coverage (i.e., $x_{i-1}(S) > x_i(S) > 0$). Define $\mathcal{Y}_i \triangleq \{S_j : x_{i-1}(S_j) > x_i(S_j)\}$, namely \mathcal{Y}_i contains the sets that lose coverage due to the invocation of **insert** during the i th iteration.

If no coverage is lost during iteration i , namely if $\mathcal{Y}_i = \emptyset$ (or $Z_i = \emptyset$), then $\sum_{j < i} \sum_{u \in U} w(u)[z_{i-1}(u, S_j) - z_i(u, S_j)] = 0$, and we are done since $x_i(S) > 0$ implies that $\rho(\vec{z}_i, \vec{x}_i, S_i) > 0$.

If $\mathcal{Y}_i \neq \emptyset$, then by Claim 7 and the code of **insert** it follows that $\rho(z_{i-1}, x_{i-1}, S) \leq \rho(z_i, x_i, S')$, for every $S \in \mathcal{Y}_i$ and S' such that $x_i(S') > 0$. In words, the efficiency of the sets that retain coverage is at least as high as the efficiency of the sets that lost coverage. As mentioned above, observe that there may be at most one set S_j , such that $S_j \in \mathcal{Y}_i$ and $x_i(S_j) > 0$. Define $\rho_{\min}^i \triangleq \min \{\rho(\vec{z}_{i-1}, \vec{x}_{i-1}, S_j) : x_i(S_j) > 0, j < i\}$ and $\rho(\mathcal{Y}_i) \triangleq \max \{\rho(\vec{z}_{i-1}, \vec{x}_{i-1}, S_j) : S_j \in \mathcal{Y}_i\}$, and we have that $\rho(\mathcal{Y}_i) \leq \rho_{\min}^i$.

Define W_i^L to be the total weight of lost coverage during the i th iteration and define W_i^R to be the total weight of retained coverage at the i th iterations. That is, define: $W_i^L \triangleq \sum_{j < i} \sum_{u \in U} w(u)[z_{i-1}(u, S_j) - z_i(u, S_j)]$ and $W_i^R \triangleq \sum_{j < i} \sum_{u \in U} w(u)z_i(u, S_j)$. Observe that $w(\vec{z}_{i-1}) = W_i^L + W_i^R$. We have that $W_i^R = \sum_{j < i} \rho(\vec{z}_{i-1}, \vec{x}_{i-1}, S_j)x_i(S_j)c(S_j) \geq \rho_{\min}^i \sum_{j < i} x_i(S_j)c(S_j)$, and

$$\begin{aligned} W_i^L &= \sum_{j < i} \rho(\vec{z}_{i-1}, \vec{x}_{i-1}, S_j)[x_{i-1}(S_j)c(S_j) - x_i(S_j)c(S_j)] \\ &= \sum_{S_j \in \mathcal{Y}_i} \rho(\vec{z}_{i-1}, \vec{x}_{i-1}, S_j)[x_{i-1}(S_j)c(S_j) - x_i(S_j)c(S_j)] \\ &\leq \rho(\mathcal{Y}_i) \sum_{S_j \in \mathcal{Y}_i} [x_{i-1}(S_j)c(S_j) - x_i(S_j)c(S_j)] \\ &= \rho(\mathcal{Y}_i) \sum_{j < i} [x_{i-1}(S_j)c(S_j) - x_i(S_j)c(S_j)] \\ &\leq \rho_{\min}^i \left[1 - \sum_{j < i} x_i(S_j)c(S_j) \right] \\ &\leq W_i^R \cdot \frac{1 - \sum_{j < i} x_i(S_j)c(S_j)}{\sum_{j < i} x_i(S_j)c(S_j)}. \end{aligned}$$

It follows that

$$w(\vec{z}_{i-1}) = W_i^L + W_i^R \geq W_i^L + W_i^L \cdot \frac{\sum_{j < i} x_i(S_j)c(S_j)}{1 - \sum_{j < i} x_i(S_j)c(S_j)} = \frac{W_i^L}{1 - \sum_{j < i} x_i(S_j)c(S_j)}.$$

Now since $\mathcal{Y}_i \neq \emptyset$, we have that **insert** decreased coverage of at least one set and by the code of **insert** this implies that $\sum_S x_i(S)c(S) = 1$. Hence we have that $w(\vec{z}_{i-1}) \geq \frac{W_i^L}{x_i(S_i)c(S_i)}$, and the lemma follows. \blacktriangleleft

We are now ready to give in the next lemma an upper bound on the amount of blue money that is allotted to any element $u \in U$ at any given time.

► **Lemma 9.** *At any given time i and for every $u \in U$, $\text{blue}_i(u) \leq w(u)\hat{z}_i(u) \cdot \frac{\alpha}{\alpha-1}$.*

Proof. In this proof we consider separately new blue money, that was not transferred yet, and old blue money, that was transferred at least once. Observe that for each element $u \in U$

there is at most one index j such that $u \in F_j$. We denote this index by $f(u)$. We prove by induction on i that at any given time i and for every $u \in U$, we have that

1. $\text{blue-old}_i(u) \leq w(u) \hat{z}_i(u) \cdot \frac{1}{\alpha-1}$, and
2. $\text{blue-new}_i(u) \leq \begin{cases} 0 & i < f(u), \\ w(u) \hat{z}_i(u) & i \geq f(u). \end{cases}$

For $i = 0$, i.e., before the first input set arrives, the claims hold as there is no (blue) money. We now prove the claims for time $i \geq 1$, assuming that the claims hold for time $i - 1$. We analyze the changes in $\text{blue-old}_i(\cdot)$ and $\text{blue-new}_i(\cdot)$ taking into account one by one, in their order, the operations of the charging scheme defined above. Recall that $\hat{z}'_i(v) \triangleq \sum_{j=1}^{i-1} z_i(v, S_j)$ and that $Z_i \triangleq \{v : \hat{z}'_i(v) < \hat{z}_{i-1}(v)\}$, and observe that the out-transfer and in-transfer phases are performed only if procedure **insert** is invoked.

- **Out-transfer of money.** Blue money (old or new) is removed from elements $u \in Z_i$. For such an element u we have at the end of the out-transfer phase that $\text{blue}_i(u) = \text{blue}_{i-1}(u) \cdot \frac{\hat{z}'_i(u)}{\hat{z}_{i-1}(u)}$. By the induction hypothesis this is at most

$$\begin{aligned} \text{blue-old}_i(u) &\leq w(u) \hat{z}_{i-1}(u) \cdot \frac{1}{\alpha-1} \cdot \frac{\hat{z}'_i(u)}{\hat{z}_{i-1}(u)} = w(u) \hat{z}'_i(u) \cdot \frac{1}{\alpha-1} \\ \text{blue-new}_i(u) &\leq w(u) \hat{z}_{i-1}(u) \cdot \frac{\hat{z}'_i(u)}{\hat{z}_{i-1}(u)} = \begin{cases} 0 & i-1 < f(u), \\ w(u) \hat{z}'_i(u) & i-1 \geq f(u). \end{cases} \end{aligned}$$

For $u \in U \setminus Z_i$, $\text{blue-old}_i(u) = \text{blue-old}_{i-1}(u)$ and $\text{blue-new}_i(u) = \text{blue-new}_{i-1}(u)$. Using the induction hypothesis, at the end of the out-transfer phase, the same bound holds for $u \in U \setminus Z_i$ as well.

- **In-transfer of money.** First notice that there is no in-transfer of new blue money. Each element $u \in S_i$ gets old blue money in the amount of

$$\frac{z_i(u, S_i) \cdot w(u)}{\sum_{v \in S_i} z_i(v, S_i) \cdot w(v)} \cdot \sum_{v \in Z_i} \left(1 - \frac{\hat{z}'_i(v)}{\hat{z}_{i-1}(v)}\right) \text{blue}_{i-1}(v).$$

By the inductive hypothesis this is at most

$$\frac{z_i(u, S_i) \cdot w(u)}{\sum_{v \in S_i} z_i(v, S_i) \cdot w(v)} \cdot \sum_{v \in Z_i} (\hat{z}_{i-1}(v) - \hat{z}'_i(v)) \cdot w(v) \cdot \frac{\alpha}{\alpha-1},$$

and by Lemma 8 it follows that this is at most

$$\frac{z_i(u, S_i) \cdot w(u)}{\sum_{v \in S_i} z_i(v, S_i) \cdot w(v)} \cdot \frac{1}{\alpha-1} \cdot \rho(\vec{z}_i, \vec{x}_i, S_i) x_i(S_i) c(S_i) \leq z_i(u, S_i) \cdot w(u) \cdot \frac{1}{\alpha-1}.$$

Hence, using the upper bound on $\text{blue-old}_i(u)$ at the end of the out-transfer phase, we have that at the end of in-transfer phase:

$$\text{blue-old}_i(u) \leq w(u) \hat{z}'_i(u) \cdot \frac{1}{\alpha-1} + w(u) z_i(u, S_i) \cdot \frac{1}{\alpha-1} = w(u) \hat{z}_i(u) \cdot \frac{1}{\alpha-1}.$$

Since no transfer of new blue money occurs by the scheme, it holds at the end of the out transfer phase, like at the end of the in-transfer phase that:

$$\text{blue-new}_i(u) \leq \begin{cases} 0 & i-1 < f(u), \\ w(u) \hat{z}'_i(u) & i-1 \geq f(u). \end{cases}$$

- **Creation of new money.** if $S_i \in \text{OPT}$, then any $u \in F_i$ receives $\hat{z}_i(u)w(u)$ new blue money. Observe that for such u , $f(u) = i$. Hence, using our claims as to the end of the in-transfer phase, the inductive claim holds at the end of the creation-of-new-money phase. ◀

We now give an upper bound on the amount of red money an element may have. The proof of the following lemma is omitted from this extended abstract.

- **Lemma 10.** *At any given time i and for every $u \in U$, $\text{red}_i(u) \leq w(u)\hat{z}_i(u) \cdot \alpha \cdot c(\text{OPT}_i)$.*

Using Lemmas 9 and 10 we now give a lower bound on the weight of the fractional solution (\vec{z}, \vec{x}) .

- **Lemma 11.** *$w(\vec{z}) \geq \frac{1}{4}w(\cup_{S \in \text{OPT}} S)$ at termination.*

Proof. First observe that the total amount of blue and red money created during the course of the run is equal to the weight of the elements covered by OPT , and that all created money remains in the system, held by the various elements $u \in U$, until the end of the run. We therefore compare the total amount of money held at the end by the elements $u \in U$, to the weight of the elements covered by the online algorithm.

Let n be the number of sets in the input sequence. Since $c(\text{OPT}) = c(\text{OPT}_n) \leq 1$, we have, for $\alpha = 2$ and using Lemma 9 and Lemma 10, that for each $u \in U$, $\text{blue}_n(u) + \text{red}_n(u) \leq w(u)\hat{z}_n(u) \cdot (\alpha + \frac{\alpha}{\alpha-1}) = w(u)\hat{z}_n(u) \cdot 4$. ◀

We now give a lower bound on the weight of the (integral) solution returned by the online algorithm, in terms of the value of the fractional solution (\vec{z}, \vec{x}) .

- **Lemma 12.** *Let ALG_i be the integral solution returned by the online algorithm after processing set S_i . Then, $w(\cup_{S \in \text{ALG}_i} S) \geq (1-r) \cdot w(\vec{z}_i)$.*

Proof. By definition we have that $\text{ALG}_i = \{S : x_i(S) = 1\}$. Each such set (fully) covers all its elements, hence $w(\cup_{S \in \text{ALG}_i} S) \geq \sum_u \sum_{S \in \text{ALG}_i} w(u)z_i(u, S)$. By Claim 7 there may be at most one set S' such that $x_i(S') \in (0, 1)$. If such a set does not exist, then $w(\vec{z}_i) = \sum_u \sum_{S \in \text{ALG}_i} w(u)z_i(u, S)$, and we are done. If there exists a set S' such that $x_i(S') \in (0, 1)$, then by Claim 7 we have that $\sum_{S \in \text{ALG}_i} c(S) = 1 - x_i(S')c(S') > 1 - c(S')$. Also, due to Claim 7, $\rho(\vec{z}_i, \vec{x}_i, S') \leq \rho(\vec{z}_i, \vec{x}_i, S)$, for any $S \in \text{ALG}_i$. Therefore

$$\begin{aligned} w(\cup_{S \in \text{ALG}_i} S) &\geq \sum_{S \in \text{ALG}_i} \sum_u w(u)z_i(u, S) \\ &\geq \frac{1 - c(S')}{c(S')} \cdot \sum_u w(u)z_i(u, S') \\ &\geq \frac{1-r}{r} \cdot \sum_u w(u)z_i(u, S'). \end{aligned}$$

(Recall that r is the highest set-cost appearing in the input sequence). It follows that $w(\cup_{S \in \text{ALG}_i} S) \geq (1-r) \cdot \sum_{S \in \text{ALG}_i \cup \{S'\}} \sum_u w(u)z_i(u, S) = (1-r) \cdot w(\vec{z}_i)$. ◀

It remains to give an upper bound on the competitive ratio of the algorithm.

- **Theorem 13.** *Algorithm 2-greedy is $\frac{4}{1-r}$ -competitive.*

Proof. By Lemma 11 we have that $w(\vec{z}) \geq \frac{1}{4}w(\cup_{S \in \text{OPT}} S)$ at termination. By Lemma 12 the total weight of the elements covered by the online algorithm is at least $(1-r) \cdot w(\vec{z})$. ◀

4 $O(1)$ -competitive Randomized Algorithm

In this section we give a randomized online algorithm with an $O(1)$ competitive ratio. The algorithm is based on the deterministic $\frac{4}{1-r}$ -competitive algorithm from Section 3.

The algorithm is a barely random algorithm that chooses to run one of the following two algorithms with probability $1/2$ each:

1. Always keep a single set S_j , which is the set with the highest weight seen so far.
2. Run algorithm α -greedy of Section 3, with $\alpha = 2$, only on sets with cost at most $\frac{1}{3}$.

► **Theorem 14.** *There is a 16-competitive randomized online algorithm for OBMC.*

Proof. Let OPT be an optimal solution, and define $\text{OPT}_> = \{S \in \text{OPT} : c(S) > \frac{1}{3}\}$ and $\text{OPT}_\leq = \{S \in \text{OPT} : c(S) \leq \frac{1}{3}\}$. Observe that $w(\cup_{S \in \text{OPT}_>} S) + w(\cup_{S \in \text{OPT}_\leq} S) \geq w(\cup_{S \in \text{OPT}} S)$ and that $|\text{OPT}_>| \leq 2$. We have that at least one of the two following options must occur: (i) $w(\cup_{S \in \text{OPT}_>} S) \geq \frac{1}{4}w(\cup_{S \in \text{OPT}} S)$, or (ii) $w(\cup_{S \in \text{OPT}_\leq} S) \geq \frac{3}{4}w(\cup_{S \in \text{OPT}} S)$.

If $w(\cup_{S \in \text{OPT}_>} S) \geq \frac{1}{4}w(\cup_{S \in \text{OPT}} S)$, then there exists a single set in $\text{OPT}_>$ whose weight is at least $\frac{1}{8}w(\cup_{S \in \text{OPT}} S)$. In this case, the algorithm that keeps a single set with the maximum-weight coverage seen so far obtains a solution with weight at least $\frac{1}{8}w(\cup_{S \in \text{OPT}} S)$.

If $w(\cup_{S \in \text{OPT}_\leq} S) \geq \frac{3}{4}w(\cup_{S \in \text{OPT}} S)$, then α -greedy, for $\alpha = 2$, which runs only on the sets with cost at most $\frac{1}{3}$ is 6-competitive, and therefore it computes a solution whose weight is at least $\frac{1}{6}w(\cup_{S \in \text{OPT}_\leq} S) \geq \frac{1}{6} \cdot \frac{3}{4}w(\cup_{S \in \text{OPT}} S) = \frac{1}{8}w(\cup_{S \in \text{OPT}} S)$.

Since the algorithm chooses with equal probability $\frac{1}{2}$ each one of the two algorithms, we have that $\mathbb{E}[\text{ALG}] \geq \frac{1}{2} \cdot \frac{1}{8} \cdot w(\cup_{S \in \text{OPT}} S) = \frac{1}{16}w(\cup_{S \in \text{OPT}} S)$. ◀

5 $O(\Delta)$ -competitive Algorithm

In this section we present an $O(\Delta)$ -competitive algorithm for OBMC. Given an OBMC instance we define a bipartite graph $G = (S, U, E)$, where $(S, u) \in E$ if and only if $u \in S$. Given a collection of sets \mathcal{S}' , let $G[\mathcal{S}']$ be the subgraph of G that is induced by \mathcal{S}' and U . The algorithm is based on computing maximum cardinality matchings between sets (\mathcal{S}' , the left side of G) and elements (U , the right side of G). Let **MaxMatch** be an algorithm that solves the MAXIMUM CARDINALITY MATCHING problem in bipartite graphs.

The OBMC algorithm works as follows. Upon arrival of a set S_i , $i \geq 1$, the algorithm constructs a solution \mathcal{S}_i using the previous solution \mathcal{S}_{i-1} (we initialize $\mathcal{S}_0 = \emptyset$). The algorithm looks for an element to be matched to S_i , where S_i takes precedence over sets that cost more than $c(S_i)$. This is done as follows. First a maximum matching is computed for the collection of sets that includes those sets in the already-computed solution that have cost at most $c(S_i)$, and S_i . If it is impossible to match all these sets, or their total cost exceeds 1, S_i is *rejected*, and the current solution remains unchanged. Otherwise S_i is *accepted*, and all the sets in the already-computed solution that have cost at most $c(S_i)$ are not dropped. In the latter case the algorithm then tries to extend the matching by assigning an element to those sets in \mathcal{S}_{i-1} that cost more than $c(S_i)$. Such a set S_j is *dropped* if a matching cannot be obtained or if the total cost exceeds 1. Algorithm 3 shows a formal pseudo-code of this algorithm.

In what follows ALG denotes the collection of sets that is output by the algorithm. We first prove that the solution is feasible.

► **Lemma 15.** $c(\text{ALG}) \leq 1$.

Proof. We prove by induction on i that after each iteration we have $c(\mathcal{S}_i) \leq 1$. The base case is trivial, since $\mathcal{S}_0 = \emptyset$. For the inductive step, assume that $c(\mathcal{S}_{i-1}) \leq 1$. \mathcal{S}_i is initially

Algorithm 3: Match; operations when set S_i arrives.

```

1  $\mathcal{S}_i \leftarrow \{S \in \mathcal{S}_{i-1} : c(S) \leq c(S_i)\}$ 
2  $M \leftarrow \text{MaxMatch}(G[\mathcal{S}_i \cup \{S_i\}])$ 
3 if  $|\mathcal{S}_i| = |\mathcal{S}_{i-1} \cup \{S_i\}|$  then
4   if  $c(\mathcal{S}_i \cup \{S_i\}) \leq 1$  then
5      $\mathcal{S}_i \leftarrow \mathcal{S}_i \cup \{S_i\}; M_i \leftarrow M$ 
6      $\mathcal{S}'_i \leftarrow \mathcal{S}_{i-1} \setminus \mathcal{S}_i$ 
7     while  $\mathcal{S}'_i \neq \emptyset$  do
8        $j \leftarrow \text{argmin}_j \{c(S_j) : S_j \in \mathcal{S}'_i\}$ 
9        $\mathcal{S}'_i \leftarrow \mathcal{S}'_i \setminus \{S_j\}$ 
10       $M \leftarrow \text{MaxMatch}(G[\mathcal{S}_i \cup \{S_j\}])$ 
11      if  $|\mathcal{S}_i| = |\mathcal{S}_i \cup \{S_j\}|$  then
12        if  $c(\mathcal{S}_i \cup \{S_j\}) \leq 1$  then
13           $\mathcal{S}_i \leftarrow \mathcal{S}_i \cup \{S_j\}; M_i \leftarrow M$ 
14    else
15       $\mathcal{S}_i \leftarrow \mathcal{S}_{i-1}; M_i \leftarrow M_{i-1}$ 
16 else
17    $\mathcal{S}_i \leftarrow \mathcal{S}_{i-1}; M_i \leftarrow M_{i-1}$ 

```

a subset of \mathcal{S}_{i-1} (Line 1) and therefore it is feasible at this stage due to the inductive hypothesis. Furthermore, we add sets to \mathcal{S}_i only if the budget constraint is not violated (Lines 4 and 12). \blacktriangleleft

We now show that the number of sets in the solution never decreases, and may increase by at most one set in any step (proof omitted from this extended abstract).

► **Lemma 16.** $|\mathcal{S}_{i-1}| \leq |\mathcal{S}_i| \leq |\mathcal{S}_{i-1}| + 1$, for every $i > 0$.

Since $|\text{ALG}| = |M_m|$ and $|M_m| \leq |\cup_{S \in \text{ALG}} S|$ we have the following observation.

► **Observation 17.** $|\cup_{S \in \text{ALG}} S| \geq |\text{ALG}|$.

Let OPT denote an optimal collection of sets. We partition $\text{OPT} \setminus \text{ALG}$ into two collections. Let τ be the cost of the cheapest set that was either rejected upon arrival due to the budget constraint, or dropped later due to the budget constraint. More formally, define

$$\begin{aligned} \tau_1 &= \min\{c(S_i) : S_i \text{ was rejected by Line 4 at the } i\text{th iteration}\} \cup \{\infty\}, \\ \tau_2 &= \min\{c(S_i) : S_i \text{ was dropped by Line 12 at the } j\text{th iteration}\} \cup \{\infty\}, \\ \tau &= \min\{\tau_1, \tau_2\}, \end{aligned}$$

and define

$$\text{OPT}' = \{S \in \text{OPT} \setminus \text{ALG} : c(S) < \tau\} \quad \text{and} \quad \text{OPT}'' = \{S \in \text{OPT} \setminus \text{ALG} : c(S) \geq \tau\}.$$

The next two lemmas essentially give the upper bound on the competitive ratio of the algorithm. Their proofs are omitted from this extended abstract. The first lemma shows that the elements that are covered by the sets in OPT' are covered by the sets in ALG . The second one shows that the size of OPT'' is bounded from above by the size of ALG .

► **Lemma 18.** $\cup_{S \in \text{OPT}'} S \subseteq \cup_{S \in \text{ALG}} S$.

► **Lemma 19.** $|\text{OPT}''| \leq |\text{ALG}|$.

It remains to give an upper bound on the competitive ratio of Algorithm **Match**.

► **Theorem 20.** *Algorithm **Match** is $(\Delta + 2)$ -competitive.*

Proof. We have that

$$\begin{aligned} w(\cup_{S \in \text{OPT}} S) &\leq w(\cup_{S \in \text{OPT} \cap \text{ALG}} S) + w(\cup_{S \in \text{OPT}'} S) + w(\cup_{S \in \text{OPT}''} S) \\ &\leq 2 \cdot w(\cup_{S \in \text{ALG}} S) + w(\cup_{S \in \text{OPT}''} S) \end{aligned} \tag{1}$$

$$\leq 2 \cdot w(\cup_{S \in \text{ALG}} S) + |\text{OPT}''| \cdot \Delta \tag{2}$$

$$\leq 2 \cdot w(\cup_{S \in \text{ALG}} S) + |\text{ALG}| \cdot \Delta \tag{3}$$

$$\leq (\Delta + 2) \cdot w(\cup_{S \in \text{ALG}} S), \tag{4}$$

where (1) follows from Lemma 18, (2) follows from the assumption that the weight of an element is at least 1, (3) is due to Lemma 19, and (4) is due to Observation 17 and the above assumption. ◀

5.1 $w(U)$ is known in advance

We obtain a deterministic algorithm by running algorithm **Match** with the additional rule that if a set S_i with $w(S_i) \geq \sqrt{w(U)}$ arrives, then we drop all currently taken sets, take set S_i , and never change further the solution (i.e., the final output solution is $\{S_i\}$).

Observe that if there exists a set $S_i \in \mathcal{S}$ with $w(S_i) \geq \sqrt{w(U)}$, then the total weight of the elements in the sets of ALG is at least $\sqrt{w(U)}$, while the total weight of the elements in the sets of OPT is at most $w(U)$. If no such set exists then ALG is equivalent to **Match** which is $(\Delta + 2) < (\sqrt{w(U)} + 2)$ -competitive. This leads to the following result.

► **Theorem 21.** *There exists a deterministic online algorithm whose competitive ratio is $O(\min\{\Delta, \sqrt{w(U)}\})$, provided that $w(U)$ is known in advance.*

Note that for the unit-weight case $w(U)$ equals n , and we thus get an $O(\min\{\Delta, \sqrt{n}\})$ -competitive deterministic algorithm for this case.

6 Conclusions

We have studied the ONLINE BUDGETED MAXIMUM COVERAGE problem. We presented a deterministic online algorithms in terms of three parameters of the given instance, and we gave deterministic lower bounds bases on these parameters. We also provided a randomized $O(1)$ -competitive algorithm. Finally, both our upper and lower bounds on the deterministic competitive ratio apply to REMOVABLE ONLINE KNAPSACK which is the preemptive version of the ONLINE KNAPSACK problem.

We briefly mention some possible future research directions. Obvious open problems are closing the gap between the deterministic upper and lower bounds ($O(\frac{1}{1-r})$ vs. $\Omega(\frac{1}{\sqrt{1-r}})$), and decreasing the (constant) randomized competitive ratio. Other interesting goals would be to design an $O(\sqrt{w(U)})$ -competitive deterministic algorithm that does not require advance knowledge of $w(U)$, and devising a single deterministic algorithm that obtains as a competitive ratio the minimum of all deterministic competitive ratios shown in this paper.

References

- 1 Alexander A. Ageev and Maxim Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3):307–328, 2004.
- 2 Noga Alon, Baruch Awerbuch, Yossi Azar, Niv Buchbinder, and Joseph Naor. The online set cover problem. *SIAM Journal on Computing*, 39(2):361–370, 2009.
- 3 Noga Alon, Dana Moshkovitz, and Shmuel Safra. Algorithmic construction of sets for k -restrictions. *ACM Transactions on Algorithms*, 2(2):153–177, 2006.
- 4 Giorgio Ausiello, Nicolas Boria, Aristotelis Giannakos, Giorgio Lucarelli, and Vangelis Th. Paschos. Online maximum k -coverage. *Discrete Applied Mathematics*, 160(13-14):1901–1913, 2012.
- 5 Baruch Awerbuch, Yossi Azar, Amos Fiat, and Frank Thomson Leighton. Making commitments in the face of uncertainty: How to pick a winner almost every time. In *28th Annual ACM Symposium on the Theory of Computing*, pages 519–530, 1996.
- 6 R. Bar-Yehuda and S. Even. A linear time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms*, 2:198–203, 1981.
- 7 Oded Berman, Dimitris Bertsimas, and Richard C. Larson. Locating discretionary service facilities, ii: Maximizing market size, minimizing inconvenience. *Operations Research*, 43(4):623–632, 1995.
- 8 Oded Berman, Richard C. Larson, and Nikoletta Fouska. Optimal location of discretionary service facilities. *Transportation Science*, 26(4):201–611, 1992.
- 9 Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. Submodular maximization with cardinality constraints. In *25th Annual ACM-SIAM Symp. on Discrete Algorithms*, pages 1433–1452, 2014.
- 10 Niv Buchbinder, Moran Feldman, and Roy Schwartz. Online submodular maximization with preemption. In *26th Annual ACM-SIAM Symp. on Discrete Algorithms*, pages 1202–1216, 2015.
- 11 Niv Buchbinder and Joseph Naor. Online primal-dual algorithms for covering and packing. *Math. Oper. Res.*, 34(2):270–286, 2009.
- 12 Amit Chakrabarti and Sagar Kale. Submodular maximization meets streaming: Matchings, matroids, and more. In *17th Intl. Conference on Integer Programming and Combinatorial Optimization*, volume 8494 of *LNCS*, pages 210–221, 2014.
- 13 V. Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
- 14 Marek Cygan, Lukasz Jeż, and Jiri Sgall. Online knapsack revisited. *Theory of Computing Systems*, 2016. To appear.
- 15 Marc Demange and Vangelis Th. Paschos. On-line vertex-covering. *Theoretical Computer Science*, 332(1–3):83–108, 2005.
- 16 I. Dinur, V. Guruswami, S. Khot, and O. Regev. A new multilayered PCP and the hardness of hypergraph vertex cover. *SIAM Journal on Computing*, 34(5):1129–1146, 2005.
- 17 Irit Dinur and Shmuel Safra. The importance of being biased. In *34th Annual ACM Symposium on the Theory of Computing*, pages 33–42, 2002.
- 18 Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- 19 M. L. Fisher, George L. Nemhauser, and Laurence A. Wolsey. Analysis of approximation algorithms for maximizing submodular set function II. *Mathematical Programming Study*, 8:73–87, 1978.
- 20 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, San Francisco, 1979.

- 21 Xin Han, Yasushi Kawase, and Kazuhisa Makino. Randomized algorithms for online knapsack problems. *Theoretical Computer Science*, 562:395–405, 2015.
- 22 Dorit S. Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on Computing*, 11(3):555–556, 1982.
- 23 Dorit S. Hochbaum, editor. *Approximation Algorithms for NP-Hard Problem*. PWS Publishing Company, 1997.
- 24 Dorit S. Hochbaum and Anu Pathria. Analysis of the greedy approach in problems of maximum k -coverage. *Naval Research Logistics*, 45(6):615–627, 1998.
- 25 Oscar H. Ibarra and Chul E. Kim. Fast approximation algorithms for the knapsack and sum of subset problems. *Journal of the ACM*, 22(4):463–468, 1975.
- 26 Kazuo Iwama and Shiro Taketomi. Removable online knapsack problems. In *29th Annual Intl. Colloquium on Automata, Languages and Programming*, volume 2380 of *LNCS*, pages 293–305, 2002.
- 27 D. S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.
- 28 Samir Khuller, Anna Moss, and Joseph Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39–45, 1999.
- 29 Dennis Komm, Richard Královic, and Tobias Mömke. On the advice complexity of the set cover problem. In *7th International Computer Science Symposium in Russia*, pages 241–252, 2012.
- 30 L. Lovász. On the ratio of optimal integral and fractional solutions. *Discrete Mathematics*, 13:383–390, 1975.
- 31 Alberto Marchetti-Spaccamela and Carlo Vercellis. Stochastic on-line knapsack problems. *Math. Program.*, 68:73–104, 1995.
- 32 N. Megiddo, E. Zemel, and S. L. Hakimi. The maximum coverage location problem. *SIAM Journal on Algebraic and Discrete Methods*, 4(2):253–261, 1983.
- 33 George L. Nemhauser and Laurence A. Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, Inc., 1988.
- 34 George L. Nemhauser, Laurence A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions I. *Mathematical Programming*, 14(1):265–294, 1978.
- 35 Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *29th Annual ACM Symposium on the Theory of Computing*, pages 475–484, 1997.
- 36 Barna Saha and Lise Getoor. On maximum coverage in the streaming model & application to multi-topic blog-watch. In *SIAM International Conference on Data Mining*, pages 697–708, 2009.
- 37 Sartaj Sahni. Approximate algorithms for the 0/1 knapsack problem. *Journal of the ACM*, 22(1):115–124, 1975.
- 38 Maxim Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1):41–43, 2004.
- 39 Ashwinkumar Badanidiyuru Varadaraja. Buyback problem - approximate matroid intersection with cancellation costs. In *38th Annual Intl. Colloquium on Automata, Languages and Programming*, volume 6755 of *LNCS*, pages 379–390, 2011.
- 40 Vijay V. Vazirani. *Approximation Algorithms*. Springer-Verlag, Berlin Heidelberg New York, 2001.
- 41 Yunhong Zhou, Deeparnab Chakrabarty, and Rajan M. Lukose. Budget constrained bidding in keyword auctions and online knapsack problems. In *4th international Workshop on Internet and Network Economics*, volume 5385 of *LNCS*, pages 566–576, 2008.