

Video Distribution Under Multiple Constraints

EXTENDED ABSTRACT

Boaz Patt-Shamir
School of Electrical Engineering
Tel Aviv University
Tel Aviv 69978, Israel
Email: boaz@eng.tau.ac.il

Dror Rawitz
Faculty of Science and Science Education
& Caesarea Rothschild Institute
University of Haifa
Haifa 31905, Israel
Email: rawitz@cri.haifa.ac.il

Abstract

We consider the optimization problem of providing a set of video streams to a set of clients, where each stream has costs in m possible measures (such as communication bandwidth, processing bandwidth etc.), and each client has its own utility function for each stream. We assume that the server has a budget cap on each of the m cost measures; each client has an upper bound on the utility that can be derived from it, and potentially also upper bounds in each of the m cost measures. The task is to choose which streams the server will provide, and out of this set, which streams each client will receive. The goal is to maximize the overall utility subject to the budget constraints. We give an efficient approximation algorithm with approximation factor of $O(m)$ with respect to the optimal possible utility for any input, assuming that clients have only a bound on their maximal utility. If, in addition, each client has at most m_c capacity constraints, then the approximation factor increases by another factor of $O(m_c \log n)$, where n is the input length. We also consider the special case of “small” streams, namely where each stream has cost of at most $O(1/\log n)$ fraction of the budget cap, in each measure. For this case we present an algorithm whose approximation ratio is $O(\log n)$.

1 Introduction

The following model is an abstraction of the way cable TV is distributed in many cases (see Figure 1). There are many available *streams* to multicast, and there are *clients* (or *users*), each with his own *utility* for each stream. A client may be an individual household, or a neighborhood video gateway, and the utility may represent the revenue generated by the client, or a measure of user satisfaction. The

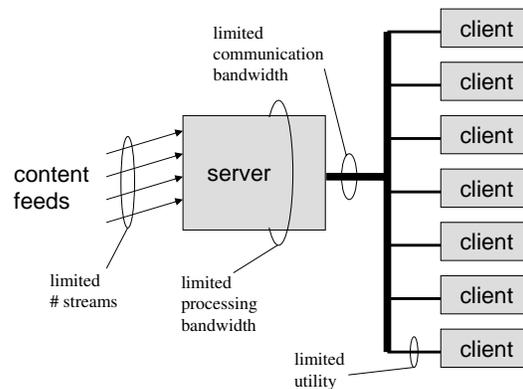


Figure 1. A schematic representation of a typical system. The server serves contents to clients, using a bounded number of input streams, bounded computational and communication bandwidth. Each client can generate bounded utility.

server (possibly a cable head-end serving video gateways, or a video gateway serving households) transmits a subset of the available streams over a multicast-capable network (typically Ethernet or DOCSIS): A transmitted stream can be received by all clients. The objective of the system is to maximize overall utility, but there are several constraints which any solution must respect. At the server, these constraints typically include limited outgoing communication bandwidth, and may also include limited processing bandwidth, limited number of input ports, etc. In general, transmitting a stream incurs a cost at the server in each of m possible measures. In our scenario, each of these m cost measures has a given *budget* cap that may not be exceeded. At the client side, the main constraint is that only a bounded amount of utility can be derived from each client. Clients

may have other constraints, like an upper bound on incoming bandwidth. In general, we assume that each client has up to m_c budgets, and each stream has a cost in each of the clients' budgets. The task is, subject to the given constraints, to select streams to broadcast by the server, and to select streams to deliver to each user, so as to maximize the overall utility of the system.

It is easy to see that finding the optimal solution to this very practical problem is computationally hard: even if there were a single user, the problem is a strict generalization of the Knapsack Problem; from another perspective, even if there were a single cost measure, and each stream had either unit cost and unit utility or zero cost and utility for each user, then the problem is a generalization of the Maximum Coverage Problem [8]. We therefore resort to near-optimal solutions, which guarantee worst-case approximation ratio with respect to the optimal solution.

Our Results. In this paper we present several algorithms for the problem. Our most general algorithm guarantees approximation factor of $O(mm_c \log n)$. If all streams have costs which do not exceed an $O(1/\log n)$ fraction of the budget, then we can guarantee $O(\log n)$ -approximation. On the other hand, if the only constraint at the client side is caps on the client utilities (even if stream costs are large), then we can guarantee to deliver at least an $\Omega(1/m)$ fraction of the best possible utility.

To state the results, we first define the problem formally. (For a complete glossary of notation, see Figure 2.)

Multi-Budget Multi-Client Distribution (MMD)

Input:

- A collection \mathcal{S} of *streams*, a set U of *users*, and two integers $m, m_c > 0$.
- A *server cost* $c_i(S) \geq 0$ for each $S \in \mathcal{S}$ and $1 \leq i \leq m$; a *user load* $k_i^u(S) \geq 0$ for each $1 \leq i \leq m_c$, stream $S \in \mathcal{S}$, and user $u \in U$.
- A *server budget* $B_i \in \mathbb{R}^+ \cup \{\infty\}$ for each $1 \leq i \leq m$, and a *user capacity* $K_i^u \in \mathbb{R}^+ \cup \{\infty\}$ for each $1 \leq i \leq m$ and user $u \in U$.
- A *user utility* $w_u(S)$ for each user u and stream S .

Output: an assignment of a set of streams $A(u)$ to each user u maximizing $\sum_{u \in U} \sum_{S \in A(u)} w_u(S)$, such that

- *Server budget constraints:* For each $1 \leq i \leq m$,

$$\sum_{S \in \cup_{u \in U} A(u)} c_i(S) \leq B_i.$$

- *User capacity constraints:* For each $1 \leq i \leq m_c$ and user u ,

$$\sum_{S \in A(u)} k_i^u(S) \leq K_i^u.$$

We also consider the special case of MMD where there is only one server budget constraint, and also there is only

Quantities related to an MMD instance:

- \mathcal{S} : streams set
- U : users set
- c_i : i th cost function
- B_i : i th budget
- m : number of server budgets
- k_i^u : i th load function of users u
- K_i^u : i th capacity of u
- m_c : number of user budgets
- $w(S) \stackrel{\text{def}}{=} \sum_{u \in \mathcal{S}} w_u(S)$: total utility of stream S
where $f(\mathcal{C}) \stackrel{\text{def}}{=} \sum_{S \in \mathcal{C}} f(S)$ for any subset $\mathcal{C} \subseteq \mathcal{S}$
and function $f : \mathcal{S} \rightarrow \mathbb{R}^+$.

Quantities related to an assignment A :

- $\mathcal{S}(A) = \bigcup_{u \in U} A(u)$, also called the *range* of A : the set of streams that are assigned to users by A .
- $c_i(A) \stackrel{\text{def}}{=} c_i(\mathcal{S}(A))$: i th cost of A
- $k_i^u(A) \stackrel{\text{def}}{=} k_i^u(A(u))$: i th load of A on u
- $w_u(A) \stackrel{\text{def}}{=} w_u(A(u))$: utility of A w.r.t. u

Figure 2. Glossary of Notation

one capacity constraint per user (i.e., $m = m_c = 1$). We refer to this special case as the **Single-Budget Multi-Client Distribution** problem, abbreviated henceforth SMD.

Before we state our results, we need to define yet another concept. Given a capacity measure i and a user u , one can compare all streams in terms of their cost-benefit ratio: how much utility is generated by a stream for unit load. We define the *local skew* of user u at capacity measure i to be the ratio between the largest and smallest cost-benefit ratios. The *local skew of an instance*, denoted α henceforth, is the maximum, over all users u and all load measures i , of the local skew of u at i . (A formal definition is given in Section 3.) Note that $\alpha \geq 1$, and equality holds iff all load functions of each user u are proportional to his utility w_u . We note that $\log \alpha = O(\log n)$ when all numbers in the input are polynomial in n (in this paper all logarithms are to base 2 unless otherwise stated).

Using the notion of local skew, we state our main result. For simplicity, we consider the case where all costs and utilities are polynomial in the input length n .

Theorem 1.1. *Consider an instance of MMD. Then*

1. An $O(mm_c \log \alpha)$ -approximation can be computed in $O(n \log n)$ time.
2. If $c_i(S) = O(\frac{B_i}{\log n})$ for all i , and $k_i^u(S) = O(\frac{K_i^u}{\log n})$ for all i, u , then an $O(\log n)$ -approximation can be found in polynomial time.

If each user has only a single budget constraint with local skew $\alpha = 1$ (which essentially means that the user is only limited by the maximal utility it can generate), then our first algorithm guarantees an $O(m)$ -approximation.

Previous work. Our model can be viewed as a generalization of the Budgeted Set Cover problem [9], which is a variant of the Set Cover problem [7]. In the set cover problem, the input consists of a collection of sets with cost for each set; the goal is to find a subcollection of sets of minimal cost, whose union is the same as the union of the complete collection. Set cover admits $O(\log n)$ approximation [13] and not better, unless $P = NP$ [5, 1].

In the budgeted set cover, the input consists of a “budget” B and a collection of sets of weighted elements, where each set has a cost. The goal is to find a subcollection of the sets whose cost is at most B , maximizing the total weight of the union. In the (unweighted) Maximum Coverage problem, the goal is to cover as many elements as possible, using at most B sets. In this case the natural greedy algorithm computes solutions whose weight is within a factor of $1 - (1 - \frac{1}{B})^B > 1 - \frac{1}{e} \approx 0.63$ from the optimum (see [10, 8]). This ratio holds even in the more general case of submodular set function maximization [11, 6]. (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, T' in the domain of f .)

Khuller, Moss and Naor [9] show that budgeted set cover can be approximated to within $\frac{e}{e-1}$, and cannot be approximated to within any smaller factor unless $NP \subseteq DTIME(n^{O(\log \log n)})$. Sviridenko [12] extends [9] to maximization of a nondecreasing submodular set function subject to a budget constraint.

Another variant is the “group budget constraint” [3], where the sets are assumed to be partitioned into disjoint “groups” and at most one set from each group may be selected to the output. The task is to maximize the size of the union of the output sets, subject to a budget constraint. [3] shows that if all sets have unit cost then approximation to within 2 is possible; if sets have different costs, the approximation factor jumps to 12. The problem we consider is a strict generalization of both variants of the budgeted set cover problem mentioned above.

The work by Awerbuch, Azar, and Plotkin [2] is also closely related to this paper. In [2] the question is whether to admit calls into a network (and how to route them), so as to maximize overall throughput subject to link capacity constraints. One important difference between the models is that in our case, the utility of a stream depends on the algorithm (which users receive the stream), whereas the “profit” of a call in [2] is part of the input.

Solution overview and paper organization. The algorithm which proves Part 1 of Theorem 1.1 applies a series of transformations (see Figure 3): First, the multi-budget

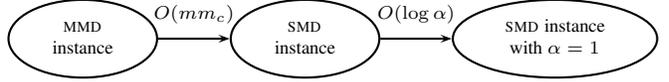


Figure 3. Description of transformations. Arrow labels indicate the approximation factors.

(MMD) instance is transformed into a single-budget (SMD) instance. Second, we show how to transform a general SMD instance into multiple SMD instances with unit skew each. Finally, we solve the SMD problem for unit skew. We describe the algorithm in a bottom-up fashion: In Section 2 we describe an $O(1)$ -approximation algorithm for SMD with unit skew, the reduction from arbitrary to unit skew is described in Section 3, and in Section 4, we describe the transformation of MMD to SMD. The algorithm for Part 2 of Theorem 1.1 is given in Section 5. This algorithm is based on ideas from [2].

2 Single Budget Constraint (SMD)

In this section we consider the case of a single budget constraint and a single capacity constraint per user with unit skew $\alpha = 1$. We give constant factor approximation algorithms for this case. Our general approach, following the work of Khuller et al. [9], is to use a greedy algorithm for this case, namely to iteratively allocate the most *cost-effective* stream to all possible users. This part is described in Section 2.1. However, the greedy algorithm is not good enough: In Section 2.2, we explain the problem and show how to fix it so as to yield a constant approximation factor.

We present an $O(n \log n)$ -time algorithm which produces utility at least $(e - 1)/2e$ times the optimal utility, if we increase the capacity of every user u by $K^u + \bar{k}^u$, where $\bar{k}^u = \max_S k^u(S)$. This is the *resource augmentation* model. Without resource augmentation, the algorithm guarantees approximation factor of $\frac{3e}{e-1}$.

Preliminaries. When the local skew is 1, the capacity and utility are at each user one and the same. Hence, in the remainder of this section, for each user u , we only consider his utility function w_u and his utility bound W_u .

In our algorithm, we may allocate a stream S to a user u even if the residual utility of the user is less than $w_u(S)$ so as to saturate the user (this happens at most once for each user). Such assignments, that satisfy the server constraints, but may violate the users’ constraints are called *semi-feasible*. We extend the definition of $w(A)$ to semi-feasible assignments as follows: $w(A) \stackrel{\text{def}}{=} \sum_u \min \{W_u, w_u(A)\}$. This means that the utility that a user u contributes is never more W_u . In a similar way we define the *fractional residual utility* of a user u for a stream S w.r.t. an assignment A

to be the utility that S adds to u if it is added to A . Formally, $\bar{w}_u(S) = 0$ for $S \in \mathcal{S}(A)$; if $S \notin \mathcal{S}(A)$, then $\bar{w}_u^A(S) = \min(w_u(S), W_u - w_u(A))$. The *fractional residual utility* of S is $\bar{w}^A(S) = \sum_u \bar{w}_u^A(S)$.

Finally, we define the *cost effectiveness* of a stream S . Given a cost function c , the cost effectiveness of S with respect to a given assignment A is defined as $\bar{w}^A(S)/c(S)$.

2.1 Basic Algorithm: Greedy

Algorithm **Greedy**, specified below, starts with the empty assignment, and iteratively adds to the solution a stream with maximum cost effectiveness with respect to the current assignment. The algorithm uses fractional residual utilities, which allows us to assign a stream S to a user u even if $\sum_{S' \in A(u)} w_u(S') > W_u - w_u(S)$. (Semi-feasible assignments are useful in the analysis, but in the final solution, the assignment is feasible.)

Algorithm 1 - Greedy($U, \mathcal{S}, c, w, W, B$)

```

1:  $A(u) \leftarrow \emptyset$ , for every  $u$ 
2:  $\mathcal{C} \leftarrow \mathcal{S}$ 
3: while  $\mathcal{C} \neq \emptyset$  do
4:   Let  $S$  be a stream that maximizes  $\bar{w}^A(S)/c(S)$ 
5:   if  $c(A) + c(S) \leq B$  then
6:      $A(u) \leftarrow A(u) \cup \{S\}$  for every  $u$  such that
        $W_u - w_u(A) > 0$ .
7:   end if
8:    $\mathcal{C} \leftarrow \mathcal{C} \setminus \{S\}$ 
9: end while
10: return  $A$ 

```

Complexity Analysis. We first consider the implementation of Algorithm **Greedy**, and explain how to get time complexity of $O(n \log n)$. The algorithm maintains a priority heap (see, e.g., [4]) of streams according to their cost effectiveness. In each iteration, we find the stream of the biggest cost effectiveness, and assign it to the users that are not yet saturated. We then delete the stream from the heap, and remove all users whose residual utility became 0. To analyze the time complexity of the algorithm, let us define a bipartite graph corresponding to the given instance of SMD: the vertex set is $\mathcal{S} \cup U$, and the edge set $E \subseteq \mathcal{S} \times U$ is defined by $(S, u) \in E$ iff $w_u(S) > 0$. Consider now an execution of the algorithm. In the iteration where stream S is the most cost effective, it is assigned to a subset U' of the users. Then we remove S from the graph, i.e., we remove S from the stream list of each user u such that $(S, u) \in E$. Also in that iteration: all users $u' \in U'$ that became saturated after the assignment of S are removed from the graph, which means that the residual utility of all streams S_1 such that $(S_1, u') \in E$ for some saturated u' needs to be updated. It follows that the total number of user updates in an iteration where S is considered is bounded by the degree of S in

the graph, and the total number of stream updates in an iteration is the sum of the degrees of all users u' that become saturated in the iteration. This means that even though the number of updates in a single iteration may be large, the total number of updates throughout the execution of the algorithm is bounded by $|E| = O(n)$.

Since the heap size is $O(|\mathcal{S}|)$, the complexity of each heap operation is $O(\log |\mathcal{S}|) = O(\log n)$. Each time we update the cost-effectiveness of a stream S we must update the heap accordingly. Since this is done no more than $|E| = O(n)$ times, we get that the total running time of the algorithm is $O(n \log n)$.

Performance Analysis. We analyze the utility of the solution computed by Algorithm **Greedy** by comparing it to the utility of any semi-feasible assignment SF (including the best such assignment).

The performance guarantee of algorithm **Greedy** follows from the observation that the utility of semi-feasible assignments is a submodular function. More precisely, let us consider an assignment just by the set of streams provided by the server. The utility of a set of streams $\mathcal{T} \subseteq \mathcal{S}$ provided by the server for a given user u is defined by

$$w_u(\mathcal{T}) = \min \left\{ W_u, \sum_{S \in \mathcal{T}} w_u(S) \right\}.$$

Note that this definition ignores the actual assignment of streams to users, but it coincides with the utility achieved by semi-feasible assignments. Thus defined, it is easy to see that for any user u , and for any two stream sets $\mathcal{T}, \mathcal{T}'$,

$$w_u(\mathcal{T}) + w_u(\mathcal{T}') \geq w_u(\mathcal{T} \cup \mathcal{T}') + w_u(\mathcal{T} \cap \mathcal{T}'),$$

i.e., the utility of semi-feasible assignment for a single user is submodular, and hence the overall utility of semi-feasible assignments is submodular as well. We can therefore apply the result of [12] to obtain a performance guarantee.

First we define some notation. Let S_i denote the i th stream considered by the algorithm, i.e., S_i is considered in the i th iteration. Let k be the number of iterations that were executed by Algorithm **Greedy** until the first stream S_{k+1} from $\mathcal{S}(\text{SF}) \setminus \mathcal{S}(A)$ is considered, but not used by A (because its addition violates the budget constraint). For $i \leq k$, let A_i denote the assignment A after the i th iteration, i.e., after considering S_i (A_0 is the empty assignment). Also, denote by A_{k+1} the (infeasible) assignment that is obtained by adding S_{k+1} to A_k . With this notation, and the observation that the utility function of semi-feasible assignments is submodular, we obtain the following result.

Lemma 2.1.

$$w(A_{k+1}) = w(A_k) + \bar{w}^{A_k}(S_{k+1}) \geq (1 - \frac{1}{e}) \cdot w(\text{SF}).$$

We note that the use of the stream S_{k+1} is essential for the analysis, as otherwise, the ratio between the optimum utility and the utility of the solution computed by

greedy may be unbounded. As an immediate corollary to Lemma 2.1, we state below the performance guarantee of Algorithm **Greedy** by comparing the output of the algorithm with an optimal solution that has a smaller budget.

Theorem 2.2. *Let A be the solution computed by Algorithm **Greedy**, and let OPT^- denote the utility of the optimal solution with reduced budget $B - c_{\max}$, where $c_{\max} = \max\{c(S) \mid S \in \mathcal{S}\}$. Then $w(A) \geq (1 - 1/e) \cdot \text{OPT}^-$.*

2.2 Fixing the Greedy Algorithm

In Theorem 2.2, the performance of the algorithm was guaranteed only after adding the stream S_{k+1} . We now show how to modify Algorithm **Greedy** to obtain approximate assignments without resource augmentation.

First, let us explain what is the weakness of the greedy algorithm. Roughly speaking, the problem with a greedy solution is that it may assign a stream S_1 with large cost-effectiveness but low absolute utility, and S_1 may block from inclusion another stream S_2 whose cost effectiveness is slightly smaller, but whose absolute utility is much larger. For example, S_2 may require the whole bandwidth budget, so even a tiny stream S_1 that was assigned will block S_2 from being assigned.

This ‘‘hole’’ in the behavior of Greedy is handled by the following trick: we find the best single-stream solution, compare it to the greedy solution, and pick the best.

More formally, let $S_{\max} = \arg\max\{w(S) \mid S \in \mathcal{S}\}$, and let A_{\max} be the assignment that assigns the single stream S_{\max} to all possible users. The modified algorithm computes assignment A_G by Algorithm **Greedy**, computes assignment A_{\max} , and outputs the better one. We denote the latter assignment by \tilde{A} . Note that \tilde{A} may still be a semi-feasible assignment. However, \tilde{A} is a $(\frac{2e}{e-1})$ -approximation:

Lemma 2.3. $w(\tilde{A}) \geq \frac{e-1}{2e} \cdot \text{OPT}$.

Proof: By Lemma 2.1, $w(A_k) + \bar{w}^{A_k}(S_{k+1}) \geq \frac{e-1}{e} \cdot \text{OPT}$. $\bar{w}^{A_k}(S_{k+1}) \leq w(S_{k+1}) \leq w(S_{\max})$ implies $w(A_k) + w(A_{\max}) \geq \frac{e-1}{e} \cdot \text{OPT}$, and the lemma follows. \square

A performance guarantee with resource augmentation follows directly:

Corollary 2.4. *There exists an algorithm that computes $(\frac{2e}{e-1})$ -approximations that may use a capacity of $K^u + \bar{k}^u$ for every user u , where $\bar{k}^u = \max\{k^u(S) \mid S \in \mathcal{S}\}$.*

We are also able to obtain an approximation algorithm that does not rely on resource augmentation. A crude lower bound can be obtained as follows.

Theorem 2.5. *There exists an $O(n \log n)$ time $(\frac{3e}{e-1})$ -approximation algorithm for the SMD problem.*

Proof: Consider the assignment A that was computed by the greedy algorithm. Define A_1 to be the assignment that picks, from each user, the stream that exceeds the user constraint (there may be at most one such stream for each user), and define A_2 to be the assignment that assigns to each user only the stream that fit completely within the user constraints (i.e., $A(u) = A_1(u) \cup A_2(u)$ for every u). Obviously, both A_1 and A_2 are feasible assignments and $w(A_1) + w(A_2) \geq w(A)$. It follows that $w(A_1) + w(A_2) + w(A_{\max}) \geq (1 - 1/e) \cdot \text{OPT}$, which means that one of A_1 , A_2 , and A_{\max} achieves approximation factor of a most $\frac{3e}{e-1}$. \square

3 Instances with Arbitrary Skew

In this section we explain how to deal with instances of SMD with arbitrary local skew. The idea is to use the ‘‘classify and select’’ approach: we reduce an instance of SMD with arbitrary skew to a set of instances of SMD where each of the new instances has $O(1)$ skew, and pick the best solution over the sub-instances.

Before we present the reduction, we formally define the local skew. Given an MMD instance, scale the k_i^u functions so that for every user u and cost measure i we have $\frac{w_u(S)}{k_i^u(S)} \geq 1$ for any stream S , with equality for at least one stream. Given this normalization, the *local skew* of the instance is defined by $\alpha \stackrel{\text{def}}{=} \max_{u,S,i} \left\{ \frac{w_u(S)}{k_i^u(S)} \right\}$. Notice that $\alpha \geq 1$ always, and equality holds iff all capacity functions of each user u are proportional to his utility w_u .

Now, suppose that we are given an SMD instance I with local skew α . We construct t SMD instances I_1, \dots, I_t , where $t = 1 + \lceil \log \alpha \rceil$. I_i is defined as follows. The streams and users are the same as in the original instance, and so are the cost function c and the budget B . We define a new utility function w_u^i for every user u :

$$w_u^i(S) = \begin{cases} k^u(S) & 2^{i-1} \leq \frac{w_u(S)}{k^u(S)} < 2^i, \\ 0 & \text{otherwise.} \end{cases}$$

That is, the i th utility function w_u^i of u only considers sets whose utility per capacity ratio is between 2^{i-1} and 2^i . We also set $W_u^i = K^u$.

Theorem 3.1. *There exists an $O(n \log n)$ time algorithm that computes $O(1 + \log \alpha)$ -approximate solutions for any instance SMD with skew α .*

Proof: Let I be a SMD instance of skew α , and let I_1, \dots, I_t be the SMD instances that are obtained as above. Clearly, each user-stream pair appear with non-zero utility in exactly one of the SMD instances I_1, \dots, I_t . Hence, $\sum_i \text{OPT}_i \geq \frac{\text{OPT}}{2}$, where OPT_i the optimum value of I_i . It

follows that there exists i such that $\text{OPT}_i \geq \frac{\text{OPT}}{2t}$. Hence, by finding an approximate solution for every SMD instance I_i , and choosing the one with maximum utility, we get an approximate solution for I .

As for the running time, let $G = (S, U, E)$ be the bipartite graph that corresponds to the problem instance I . The reduction places each edge from E in exactly one of the instances I_1, \dots, I_t . Hence, $\sum_i n_i = O(n)$, where n_i is the size of the instance I_i . By Theorem 2.5, an $O(1)$ -approximation can be computed in $O(n_i \log n_i)$ for every SMD instance I_i . It follows that the total running time is $O(\sum_i n_i \log n_i) = O(\sum_i n_i \log n) = O(n \log n)$. \square

4 Multiple Budget Constraints

In this section we show how to reduce MMD to SMD. If the server has m finite budget constraints, and a user has at most m_c budget constraints, then the reduction results in losing an approximation factor of $O(mm_c)$. The local skew may also increase by a factor of at most m_c . Our technique can be used to extend [12] to multiple budget constraints.

The main idea in the reduction is to normalize and add all cost measures to single cost, and similarly to normalize and add all capacity measures to single capacity for every user. Specifically, given an instance I_M of MMD, we construct an instance I_S of SMD as follows. The users, streams, and utility functions in I_S are just the same as in I_M . The single server cost function in I_S is defined by $c(S) = \sum_{i=1}^m \frac{c_i(S)}{B_i}$ for each stream $S \in \mathcal{S}$, and the single budget in I_S is $B = m$. Similarly, we define in I_S the single capacity constraint of each user u by $k^u(S) = \sum_i \frac{k_i^u(S)}{K_i^u}$ and $K^u = m_c$. This concludes the description of the input transformation. The output transformation is described later.

We first bound the skew of transformed instance.

Lemma 4.1. *Let α_S and α_M denote the skews of I_S and I_M , respectively. Then $\alpha_S \leq m_c \cdot \alpha_M$.*

Proof: We compute the local skew of I_S . First,

$$\begin{aligned} \frac{w_u(S)}{k^u(S)} &= \frac{w_u(S)}{\sum_i k_i^u(S)/K_i^u} \\ &\leq \frac{w_u(S)}{k_i^u(S)/K_i^u} = \frac{K_i^u \cdot w_u(S)}{k_i^u(S)} \leq K_i^u \cdot \alpha_M \end{aligned}$$

for every user and every i . Hence, $\frac{w_u(S)}{k^u(S)} \leq K_{\min}^u \cdot \alpha_M$, where $K_{\min}^u = \min_i K_i^u$. On the other hand,

$$\frac{w_u(S)}{k^u(S)} = \frac{w_u(S)}{\sum_i k_i^u(S)/K_i^u} \geq \frac{K_{\min}^u \cdot w_u(S)}{\sum_i k_i^u(S)} \geq \frac{K_{\min}^u}{m_c}.$$

It follows that the local skew of I_S is at most $m_c \cdot \alpha_M$. \square

Next we relate a solution to I_S to a solution to I_M .

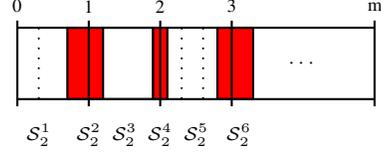


Figure 4. Decomposition of \mathcal{S}_2 . Shaded areas represent $\mathcal{S}_2^{2\ell}$ sets, and white areas represent $\mathcal{S}_2^{2\ell-1}$ sets. The dotted lines are boundaries between streams that belong to the same subset.

Lemma 4.2. *Let A be an r -approximate assignment to I_S . Then (1) $c_i(A) \leq m \cdot B_i$ for every i , (2) $k_i^u(A) \leq m_c \cdot K_i^u$ for every u and i , and (3) $w(A) \geq \text{OPT}_M/r$, where OPT_M is the optimum for I_M .*

Proof: To prove 1 and 2, note that the cost of a stream S is $c(S) = \sum_{i=1}^m \frac{c_i(S)}{B_i}$, therefore $\frac{c_i(A)}{B_i} \leq c(A) \leq B = m$. It follows that $c_i(A) \leq m \cdot B_i$ for every i . Similarly, $k_i^u(A) \leq m_c \cdot K_i^u$ for every u and i . We now prove 3. Let A^* be an optimal solution for I_M . We claim that A^* is a feasible assignment to I_S . First,

$$c(A^*) = \sum_{S \in \mathcal{S}(A^*)} \sum_{i=1}^m \frac{c_i(S)}{B_i} = \sum_{i=1}^m \frac{c_i(A^*)}{B_i} \leq \sum_{i=1}^m \frac{B_i}{B_i} = m.$$

Similarly, for every user u ,

$$\begin{aligned} k^u(A^*) &= \sum_{S \in A^*(u)} \sum_{i=1}^{m_c} \frac{k_i^u(S)}{K_i^u} = \sum_{i=1}^{m_c} \frac{k_i^u(A^*)}{K_i^u} \\ &\leq \sum_{i=1}^{m_c} \frac{K_i^u}{K_i^u} = m_c. \end{aligned}$$

Hence, $w(A^*) = \text{OPT}_M \leq \text{OPT}_S$, where OPT_S is the optimum of I_S . If A is an r -approximation for I_S , then $w(A) \geq \text{OPT}_S/r \geq \text{OPT}_M/r$, and we are done. \square

Output transformation. We now explain how to transform a solution A for I_S into a feasible solution for the original I_M . Let A be an assignment for I_S . Divide $\mathcal{S}(A)$ into two sets: \mathcal{S}_1 contains all streams whose (single) cost is larger than 1, and \mathcal{S}_2 contains all other streams. Each stream in \mathcal{S}_1 is a possible complete solution: such assignment is feasible since $c_i(S) \leq B_i$ for every S . Formally, for each $S \in \mathcal{S}_1$ we define the assignment $A|_{\{S\}}$, where $A|_{\mathcal{C}}(u) = A(u) \cap \mathcal{C}$. These are the assignments we consider from \mathcal{S}_1 . Note that $\sum_{S \in \mathcal{S}_1} c(S) \geq |\mathcal{S}_1|$, and therefore $\sum_{S \in \mathcal{S}_2} c(S) \leq m - |\mathcal{S}_1|$.

To define the assignments based on \mathcal{S}_2 , divide \mathcal{S}_2 into at most $2(m - |\mathcal{S}_1|) - 1$ sets \mathcal{S}_2^i as follows (see Figure 4).

Let each set $\mathcal{S}_j \in \mathcal{S}_2$ be represented by an interval of length $c(\mathcal{S}_j)$, and order these intervals consecutively along

the real line starting from 0, according to some arbitrary order. Now consider the integer points. For each such point $1 \leq \ell \leq m - 1$, there may be at most one stream whose interval contains each integer ℓ ; this stream (if exists) constitute the set \mathcal{S}_2^{ℓ} . The streams that lie to the right of $\ell - 1$ and to the left of ℓ constitute $\mathcal{S}_2^{\ell-1}$.

Given these $2m - 1$ subsets of $\mathcal{S}_1 \cup \mathcal{S}_2$, let A_i be the restriction of the SMD assignment to the set with largest utility. By construction, A_i satisfies the server constraints (as we show), but not necessarily the user constraints. To satisfy the user constraints, we use the same approach again. Namely, for every user u , we decompose the set $A_i(u)$ into at most $2m_c - 1$ subsets that satisfy the user capacity constraints, and remove from A_i the streams that do not belong to the subset of $A_i(u)$ of maximum utility. This completes the specification of the output transformation.

We summarize in the following theorem.

Theorem 4.3. *An r -approximation algorithm for SMD implies an $O(mm_c r)$ -approximation algorithm for MMD.*

Proof: Let A be an r -approximation for I_S , and consider the transformed output. We first argue that the output is feasible. At the server's side, if the solution is from \mathcal{S}_1 then it is feasible being a single stream, and if the solution is from \mathcal{S}_2 then it is feasible because its single cost is at most 1, and therefore its normalized cost in any measure is at most 1. Similarly, no user capacity constraint is violated.

Regarding approximation, note that the number of assignments we consider is bounded by $|\mathcal{S}_1| + 2(m - |\mathcal{S}_1|) - 1 = 2m - 1$. Hence the assignment A_i we choose has utility which is at least a $\frac{1}{2m-1}$ fraction of the utility in the solution to I_S . In the last stage, we discard streams from users to obtain assignments that adhere to user constraints, and by the same argument, we get from each user at least a $\frac{1}{2m_c-1}$ fraction of the remaining utility. The theorem follows. \square

The analysis of Theorem 4.3 is tight (details omitted). Theorem 4.3 leads us to the following result:

Theorem 4.4. *There exists an $O(n \log n)$ time $O(mm_c \log(2\alpha m_c))$ -approximation algorithm for MMD, where α is the local skew of the instance, m is the number of cost measures, m_c is the maximal number of capacity constrains at a user, and n is the input length.*

Note that if each user has only a single budget constraint with local skew $\alpha = 1$ (which means that the user is only limited by the maximal utility it can generate), then our algorithm guarantees an $O(m)$ approximation. Note further that if all costs and utilities are polynomial in the input length n , then the approximation ratio is $O(mm_c \log n)$.

As a final remark for this section, we note that our algorithm can be used to maximize arbitrary submodular set functions under m budget constraints, obtaining an $O(m)$ approximation ratio.

5 Allocating Small Streams

In this section we present an approximation algorithm for small streams. Specifically, assuming that all numbers in the input are polynomial in n , then the algorithm provides $O(\log n)$ -approximate placement for the case where each stream has cost which is at most a $O(1/\log n)$ fraction of each budget, and at most $O(1/\log n)$ fraction of each capacity. Our algorithm is based on the work of Awerbuch, Azar, and Plotkin [2].

We focus on the special case of MMD where $m_c = 1$. The extension to the case of $m_c > 1$ is straightforward.

For the sake of brevity, we assume that for every user capacity function k^u , there exists a virtual cost function c_u such that $c_u(S) = k^u(S)$ for every S , and a virtual budget $B_u = K^u$. We denote the original set of budgets by M and we abuse notation by treating U as a set of users and also as a set of budgets.

We first generalize the "local" skew α as follows. Given an MMD instance, normalize the costs such that

$$1 \leq \frac{1}{m + |U|} \cdot \frac{\sum_{u \in X} w_u(S)}{c_i(S)} \leq \gamma, \quad (1)$$

for any stream $S \in \mathcal{S}$, user set $X \subseteq U$, and cost function $i \in M \cup U$, where γ is as small as possible. The upper bound γ is called the *global skew* of the instance. The global skew γ bounds the ratio between the best and the worst streams, in terms of utility for each unit cost. Note that $\gamma \geq \alpha$ for all instances of MMD. Finally, we define $\mu \stackrel{\text{def}}{=} 2\gamma(m + |U|) + 2$.

Given an assignment A , the *normalized load* on budget i incurred by A is $L_A(i) \stackrel{\text{def}}{=} \frac{1}{B_i} \sum_{S \in \mathcal{S}(A)} c_i(S)$, and similarly, for $u \in U$, the *normalized load* is $L_A(u) \stackrel{\text{def}}{=} \frac{1}{B_u} \sum_{S \in A(u)} c_u(S)$. We also define the *exponential cost function* of budget i by $C_A(i) \stackrel{\text{def}}{=} B_i(\mu^{L_A(i)} - 1)$.

Let S_1, \dots, S_n be an arbitrary order of the streams. Algorithm **Allocate**, given formally below, starts with the empty assignment $A_0(u) = \emptyset$ for every u . Then for every stream S_j , it decides whether to allocate it and to which users, according to the exponential cost functions. Note that the maximal subset U_j may be obtained by starting with U and removing clients in decreasing order of $\frac{c_u(S_j)}{B_u} \cdot C_{A_{j-1}}(u)/w_u(S_j)$.

We start our analysis by showing that the algorithm computes feasible assignments.

Lemma 5.1. *If $c_i(S) \leq \frac{B_i}{\log \mu}$ for all i, S , then no budget constraints are ever violated.*

Proof: By contradiction. Let S_j be the first stream that caused the relative load on some budget i to exceed 1. This

Algorithm 2 - Allocate(U, \mathcal{S}, c, B, w)

1: $A_0(u) = \emptyset$ for every u
2: **for** $j = 1$ to n **do**
3: Let $C_{A_{j-1}}(i) = B_i[\mu^{L_{A_{j-1}}(i)} - 1]$ for every $i \in M \cup U$.
4: **if** there exists a maximal (inclusion wise) subset $\emptyset \subsetneq U_j \subseteq U$ such that $\sum_{i \in M \cup U_j} \frac{c_i(S_j)}{B_i} \cdot C_{A_{j-1}}(i) \leq \sum_{u \in U_j} w_u(S_j)$
 then
5: Assign S_j to the users in U_j : if $u \in U_j$ then $A_j(u) = A_{j-1}(u) \cup \{S_j\}$; otherwise $A_j(u) = A_{j-1}(u)$.
6: **else**
7: $A_j = A_{j-1}$
8: **end if**
9: **end for**

means that $L_{A_j(i)} > 1 - \frac{c_i(S_j)}{B_i}$. Since $c_i(S) \leq \frac{B_i}{\log \mu}$ by assumption, it follows that

$$\begin{aligned} \frac{C_{A_j(i)}}{B_i} &= \mu^{L_{A_j(i)}} - 1 \\ &> \mu^{1 - \frac{1}{\log \mu}} - 1 = \frac{\mu}{2} - 1 = \gamma(m + |U|). \end{aligned}$$

Hence, by the RHS of (1) we get that $\frac{c_i(S_j)}{B_i} \cdot C_{A_j}(i) > \gamma(m + |U|) \cdot c_i(S_j) \geq \sum_{u \in U_j} w_u(S_j)$ which means that stream S_j could not have been assigned to U_j . \square

We show that the approximation ratio of the algorithm is $O(1 + 2 \log \mu)$ if $c_i(S) \leq \frac{B_i}{\log \mu}$ for every stream S and $i \in M \cup U$. Let $C_j = \sum_{i \in M \cup U} C_{A_j}(i)$. Below we first show that the utility gained by the algorithm is an $\Omega(\frac{1}{\log \mu})$ fraction of C_n , and then we show that the additional utility gained by any assignment is at most C_n .

The proofs of the following two lemmas are omitted from this extended abstract.

Lemma 5.2. *Let A be the assignment that is computed by the algorithm. Then $C_n \leq 2 \log \mu \cdot w(A)$.*

Lemma 5.3. *Let A^* be an optimal assignment. Then $w(A^*) - w(A) \leq C_n$.*

Theorem 5.4. *Algorithm **Allocate** computes $O(1 + 2 \log \mu)$ -approximate solutions in the case where $c_i(S) \leq \frac{B_i}{\log \mu}$ for every stream S and cost measure i .*

Proof: By the previous two lemmas it follows that $w(A^*) - w(A) \leq 2 \log \mu \cdot w(A)$. Hence, $w(A^*) \leq (1 + 2 \log \mu) \cdot w(A)$. \square

If all numbers in the input are polynomial in n , then γ is polynomial in n , and the approximation ratio is $O(\log n)$.

Acknowledgment. Research supported in part by the Next Generation Video (NeGeV) consortium, Israel, and by the Israel Science Foundation (grant 664/05). The authors thank Baruch Awerbuch for very useful discussions.

References

- [1] N. Alon, D. Moshkovitz, and S. Safra. Algorithmic construction of sets for k -restrictions. *ACM Transactions on Algorithms*, 2(2):153–177, 2006.
- [2] B. Awerbuch, Y. Azar, and S. Plotkin. Throughput-competitive on-line routing. In *34th IEEE Symp. on Foundations of Computer Science*, pages 32–40, 1993.
- [3] C. Chekuri and A. Kumar. Maximum coverage problem with group budget constraints and applications. In *7th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, volume 3122 of LNCS, pages 72–83, 2004.
- [4] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2nd edition, 2001.
- [5] U. Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
- [6] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. Analysis of approximation algorithms for maximizing submodular set function II. *Mathematical Programming Study*, 8:73–87, 1978.
- [7] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
- [8] D. S. Hochbaum, editor. *Approximation Algorithms for NP-Hard Problem*. PWS Publishing Company, 1997.
- [9] S. Khuller, A. Moss, and J. Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39–45, 1999.
- [10] G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, 1988.
- [11] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions I. *Mathematical Programming*, 14(1):265–294, 1978.
- [12] M. Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1):41–43, 2004.
- [13] V. V. Vazirani. *Approximation Algorithms*. Springer-Verlag, Berlin Heidelberg New York, 2001.