

New Models and Algorithms for Throughput Maximization in Broadcast Scheduling

[Extended Abstract]

Chandra Chekuri^{1*}, Avigdor Gal², Sungjin Im^{1**}, Samir Khuller^{3***}, Jian Li³,
Richard McCutchen^{3 †}, Benjamin Moseley^{1‡}, and Louiqa Raschid⁴

¹ Dept. of Computer Science, Univ. of Illinois, 201 N. Goodwin Ave., Urbana, IL 61801

² Faculty of Industrial Engineering & Management, Technion – Israel Institute of Technology, Technion City, Haifa 32000, Israel.

³ Dept. of Computer Science and UMIACS, Univ. of Maryland, College Park, MD 20742.

⁴ Robert H. Smith School of Business and UMIACS, Univ. of Maryland, College Park, MD 20742.

Abstract. In this paper we consider some basic scheduling questions motivated by query processing that involve accessing resources (such as sensors) to gather data. Clients issue requests for data from resources and the data may be dynamic or changing which imposes temporal constraints on the delivery of the data. A proxy server has to compute a probing schedule for the resources since it can probe a limited number of resources at each time step. Due to overlapping client requests, multiple queries can be answered by probing the resource at a certain time. This leads to problems related to some well-studied broadcast scheduling problems. However, the specific requirements of the applications motivate some generalizations and variants of previously studied metrics for broadcast scheduling. We consider both online and offline versions of these problems and provide new algorithms and results.

1 Introduction

There is an explosion in the amount of data being produced, collected, disseminated and consumed. One important source of this data is from sensors that are being widely deployed to monitor and collect a variety of information, for example, weather and traffic. Another important source of data comes from individuals and entities publishing content on the web. Two aspects of the above type of data are the following. First, a consumer/client is typically interested only in some small subset of the available data that is relevant to her. Second, the data has temporal relevance and a client is also typically interested in data that is within some time interval of interest to her; for example traffic on a particular road during her commute window. These trends have necessitated a significant increase in the sophistication of data delivery capabilities to keep up with quantity of the data, and the need for client customization [27]. There is a large effort

* Partially supported by NSF grants CCF-0728782 and CNS-0721899.

** Supported mainly by a Samsung Fellowship, and partially by NSF grants CCF-0728782 and CNS-0721899.

*** Supported by NSF Awards CCF-0728839 and CCF-0937865, and a Google Research Award.

† Research supported by REU Award supplement to CCF-0728839

‡ Partially supported by NSF grant CNS-0721899.

in several areas of computer science to address these issues. Typically, software known as middleware handles the interface between the clients and data sources. In this paper we consider certain scheduling problems that arise in processing the queries.

Middleware primarily consists of proxy servers that collect client queries and access data sources (such as sensors) to answer queries [11,15,5,25,12,28]. In this work we consider a basic and central question that arises when the queries are time sensitive (they also may be periodic) such as “Give me the reading of sensor A at 15 mins after the hour, every hour”. The main challenge is to schedule probes to the data sources (e.g., sensors), to obtain the data at the desired time for the clients. Due to processing limitations, the proxy server is limited to probing only a small number of sensors at each time step (we assume for simplicity that it probes one sensor at each time step). However, by probing a sensor at a particular time, multiple overlapping queries requesting data from this sensor, can be answered.

More formally, there are clients that issue queries for data from a resource at a specific time, by specifying an interval of time when the resource should be queried. A central server collects all the queries and needs to design a schedule to probe the resources to answer client queries. When a resource is probed, several client queries can be answered. Typically, the queries are simple and so the computational requirements are minimal; hence we focus on the design of the probing schedule. For example, by identifying overlapping queries to the same resource we may be able to significantly reduce the number of times we query the sensors, since we can “piggyback” all the queries [26,27,28]. This overlapping nature of query processing, is very similar to the manner in which broadcast scheduling problems are approached [2,20,21,17,6]. However, the sensor probing application gives rise to new and interesting variants of broadcast scheduling problems. In this work we focus on a collection of online and offline problems motivated by the above application.

In the broadcast scheduling literature, three objectives have been the focus of study: (i) minimizing average response time⁵ [20,1,13,14,16,6,18] (and many others), (ii) minimizing maximum response time [2,6,8], and (iii) maximizing throughput [21,4,10,29]. By response time of a request, we mean the time from the arrival of the request to when it is satisfied. The first two metrics apply to settings in which all requests are to be satisfied. The third metric is relevant in situations where requests may not be satisfied beyond a certain time; in particular, the following model has been studied: Each request has a release time and deadline and it can only be satisfied within its time window and the goal is to maximize the number/weight of satisfied requests. We next explain why these metrics are not directly suitable for our purposes.

In sensor probing, the requests are time sensitive which calls for a more nuanced view of “satisfying” a request. For example, if a client requests the temperature reading, or traffic conditions at 5:30pm, then we may satisfy this query by reporting the value at 5:33pm, this would have a latency of 3 minutes. Suppose we report the value at 5:40pm with a latency of 10 minutes; the data may still be useful to the client but perhaps less than reporting the value at 5:33pm. Finally, the data may be irrelevant if the latency is say more than 20 minutes. This example demonstrates the two objectives of interest in a schedule. We are interested in “completeness”, the number of client requests that can be

⁵ Response time is also commonly referred to as flow time.

satisfied before their deadline. We are also interested in the “latency” of those requests that we do satisfy. As can be seen, previous metrics do not capture the combination of these metrics; minimizing average response time ignores deadlines and maximizing throughput with deadlines ignores the latency of satisfied requests.

In this paper we take two approaches to finding schedules that address both completeness and latency. In the first approach, we associate an arbitrary *time-dependent profit* function with each query. The profit function can take into account the impact of the latency on the value to the client. The goal then would be to find a schedule that maximizes the total profit of the requests. This model captures the previously studied maximum throughput metric, but allows more control over the quality of the schedule for queries. We consider both offline and online settings and obtain several new results. In the second approach, we directly address the tradeoff between completeness and latency. In addition to satisfying as many requests as possible we hope to also satisfy them close to when the request arrived (the arrival time could be the ideal time when the sensor should be probed). We formalize this in the following way: among all schedules that achieve a desired level of completeness, choose the one that minimizes latency of satisfied requests.

Finally, we consider another variant of the maximum throughput problem that is relevant to our application domain. In some cases, it is perfectly reasonable to report the value “before” the arrival time of the request. In the same example above, the proxy server may have the value of a sensor measuring temperature that has been probed at 5:28pm while a request for the same sensor arrives at 5:30pm. The server can use the reading at 5:28pm to answer the query. We model this aspect in two ways: by relaxing the time window of interest to the client both forwards and backwards in time, and by considering unimodal profit functions.

All the problems we consider are NP-hard in the offline setting via simple reductions from known results on broadcast scheduling [6]. We, therefore, focus on the design of efficient approximation algorithms. We also consider online variants and use the standard competitive analysis framework; for some variants we analyse the algorithms in the resource augmentation framework [19] wherein the algorithm is given extra speed over the adversary. We give below a formal description of the problems considered in the paper, followed by our results.

Problem Definitions: For convenience we shall use the standard broadcast scheduling notation of referring to pages instead of referring to sensors. We are given a set of pages $\mathcal{P} = \{p_1, \dots, p_n\}$. We assume that time is slotted, $\mathcal{T} = \{1, 2, \dots, T\}$ where T is an upper bound on the schedule length. Suppose a client sends a request for page p , which arrives at time a and is associated with deadline d . If the server broadcasts p at some time slot t such that $a \leq t \leq d$, we say the request is *satisfied*. We assume the server can broadcast at most one page in a single time slot. We use $J_{p,i}$ to denote the i th request for page p , which has the arrival time $a_{p,i} \in \mathbb{Z}^+$ and the deadline $d_{p,i} \in \mathbb{Z}^+$. Sometimes, we will consider a generalized request which may be associated with more than one interval. We use $\mathcal{T}_{p,i}$ to denote the set of time slots during which $J_{p,i}$ can be satisfied. For example, if $J_{p,i}$ has only one interval, then $\mathcal{T}_{p,i} = \{a_{p,i}, a_{p,i} + 1, \dots, d_{p,i}\}$. In this paper, we study the following objective functions.

1. **Maximizing throughput (MAX-THP):** The objective is to maximize the total number of satisfied requests. In the weighted version of MAX-THP, each request $J_{p,i}$ has a weight $w_{p,i}$. In this case, the objective is to maximize the total weight of all satisfied requests.
2. **Maximizing total profit (MAX-PFT):** This is a significant generalization of MAX-THP. In a MAX-PFT instance, each request $J_{p,i}$ is associated with an arbitrary non-negative profit function $g_{p,i} : \mathcal{T} \rightarrow \mathbb{Z}^+$. The interpretation is that the request $J_{p,i}$ obtains a value/profit of $g_{p,i}(t)$ if it is satisfied by the broadcast of p at time t . However, p may be broadcast multiple times during a schedule. In that case the request $J_{p,i}$ obtains a profit $\max_{t \in \mathcal{T}_p^A} g_{p,i}(t)$ where \mathcal{T}_p^A is the set of time slots in which p was broadcast by a given schedule. The objective is to find a schedule A such that the total profit is maximized.
3. **Completeness-Latency tradeoff:** We are given a completeness threshold $C \in (0, 1]$. The goal is to find a schedule that completes C fraction of the requests before their deadline and subject to that constraint, minimizes the latency of the completed requests.

Outline of Results: We obtain several results for the problems described above. We give a high-level description of these results below.

Maximizing Throughput and Profit: Recall that MAX-PFT is a significant generalization of MAX-THP. There is a $3/4$ -approximation for the MAX-THP problem [17] via a natural LP relaxation. We adapt the ideas in [17] to obtain a $3/4$ -approximation for the special case of MAX-PFT when the profit functions for each query are unimodal (see Section 3), which is of particular interest to our setting. Second, for the general MAX-PFT problem we obtain a $(1 - 1/e)$ -approximation, again via the natural LP relaxation. In addition, we show that the MAX-PFT problem can be cast as a special case of submodular function maximization subject to a matroid constraint. This allows us to not only obtain a different $(1 - 1/e)$ -approximation but also several generalizations and additional properties via results in [3,9]. The connection also allows us to easily show that the greedy algorithm gives a $1/2$ approximation for MAX-PFT in the *online* setting, generalizing prior work that showed this for MAX-THP [21].

We also consider how the approximation ratios and competitive ratios for MAX-THP and MAX-PFT can be improved via resource augmentation and other relaxations. We show that there is a 2-speed 1-approximation for MAX-THP. Previously, such a result was known only if all requests could be scheduled in a fractional solution [6]. In the online-setting we show that the simple greedy algorithm with s -speed achieves a $s/(s+1)$ competitive ratio for MAX-PFT. In a different direction we consider relaxing the time window in MAX-THP and prove the following result. If there is a fractional schedule that satisfies all the client requests (obtained by solving the LP relaxation to the IP), then there is an integral schedule with the following property: each request $J_{p,i}$ is satisfied in a window $[a_{p,i} - L, d_{p,i} + L]$ where $L = d_{p,i} - a_{p,i}$ is the window length of $J_{p,i}$. In other words, by either left shifting the window or right shifting the window by its length, we can always satisfy the request.

Completion-Latency Tradeoff: We show that there is an interesting tradeoff that can be obtained between latency and completeness when each request has an associated

deadline. Given a fractional LP solution (obtained by relaxing the IP) for minimizing the total latency subject to a certain completeness level, we show that we can use randomized rounding to obtain a schedule with the following properties: the expected completeness of the schedule is $\frac{3}{4}C$, where C is the completeness of the fractional schedule and the expected latency of the scheduled requests is $D(C)$ where $D(C)$ is the minimum fractional latency with completeness requirement C . The details of this result are deferred to a full version of the paper.

We also prove another result in broadcast scheduling that is of interest. This concerns the problem of minimizing the maximum response time. The first-in-first-out (FIFO) algorithm is 2-competitive in the online setting [2,6,8] and this is also the best known off-line approximation known. Moreover, it is known that in the online setting no deterministic algorithm is $(2 - \epsilon)$ -competitive for any $\epsilon > 0$ [2,6]. Here, we show that the same lower bound holds even for *randomized* online algorithms in the oblivious adversary model. The details of this result are omitted due to space constraints.

2 Preliminaries

Several of our results rely on the dependent randomized rounding framework of [17]. We first describe the LP relaxation for MAX-THP that is used as the basis for the rounding process.

2.1 An LP Relaxation for MAX-THP

We consider a natural integer programming (IP) formulation for MAX-THP. We use the indicator variable $Y_p^{(t)}$. $Y_p^{(t)} = 1$ if page p is broadcast in time-slot t and $Y_p^{(t)} = 0$ otherwise. In addition we define variables $X_{p,i}$ for the request $J_{p,i}$. This variable is 1 if and only if $J_{p,i}$ is satisfied.

$$\begin{aligned}
& \text{maximize} && \sum_{p,i} w_{p,i} X_{p,i} && (1) \\
& \text{subject to} && \forall p, t, \sum_{t \in \mathcal{T}_{p,i}} Y_p^{(t)} \geq X_{p,i} && \text{(If } p \text{ is not broadcast in } \mathcal{T}_{p,i}, J_{p,i} \text{ cannot be satisfied)} \\
& && \forall t, \sum_p Y_p^{(t)} \leq 1 && \text{(One page broadcast at one time-slot)} \\
& && \forall p, t, X_{p,i} \in \{0, 1\}, Y_p^{(t)} \in \{0, 1\}
\end{aligned}$$

By letting the domain of $X_{p,i}$ and $Y_p^{(t)}$ be $[0, 1]$, we obtain the linear programming (LP) relaxation for the problem.

2.2 Dependent rounding scheme of [17]

We briefly describe the dependent randomized rounding method of [17]. Suppose we are given a bipartite graph (A, B, E) with bipartition (A, B) . We are also given a value $x_{i,j} \in [0, 1]$ for each edge $(i, j) \in E$. The scheme in [17] provides a randomized polynomial-time algorithm that rounds each $x_{i,j}$ to a random variable $X_{i,j} \in \{0, 1\}$,

effectively keeping or dropping the edge, in such a way that the following properties hold.

(P1): Marginal distribution. For every edge (i, j) , $\Pr[X_{i,j} = 1] = x_{i,j}$.

(P2): Degree-preservation. Consider any vertex $i \in A \cup B$. Define its fractional degree d_i to be $\sum_{j:(i,j) \in E} x_{i,j}$, and integral degree D_i to be the random variable $\sum_{j:(i,j) \in E} X_{i,j}$. Then, $D_i \in \{\lfloor d_i \rfloor, \lceil d_i \rceil\}$. Note in particular that if d_i is an integer, then $D_i = d_i$ with probability 1.

(P3): Negative correlation. For any vertex i and any subset S of the neighbors of i :

$$\forall b \in \{0, 1\}, \Pr\left[\bigwedge_{j \in S} (X_{i,j} = b)\right] \leq \prod_{j \in S} \Pr[X_{i,j} = b]. \quad (2)$$

We refer the reader to [17] for more details.

3 Throughput and Profit Maximization

This section is devoted to offline and online algorithms for MAX-THP and MAX-PFT.

3.1 Offline Algorithms

Maximizing the Total Profit In this section, we consider the profit maximization (MAX-PFT) problem. Recall that in a MAX-PFT instance, each request $J_{p,i}$ is associated with a profit function $g_{p,i}(t) \geq 0$ that is an arbitrary non-negative function of the time it is satisfied. If a request for page p is satisfied multiple times by a schedule A , the profit we can get for p is the maximum one, i.e., $\max_{t \in \mathcal{T}_p^A} g_{p,i}(t)$. The objective is to find a schedule A such that the total profit is maximized. Note that MAX-THP is just a special case of MAX-PFT where the profit function $g_{p,i}(t)$ is 1 for $a_{p,i} \leq t \leq d_{p,i}$.

First, we show how to reduce MAX-PFT to MAX-THP with weighted requests where each request may have multiple intervals. We use a simple slicing trick described as follows. Consider a single request $J_{p,i}$, and let $v_1 < v_2 < \dots < v_r$ be the distinct nonnegative values taken on by its profit function $g_{p,i}$. Let $v_0 = 0$. We create r new requests for the throughput maximization instance, say $J_{p,i,j}$, $1 \leq j \leq r$, which all require page p . $J_{p,i,j}$ has weight $v_j - v_{j-1}$ and intervals consisting of time slots $\{t \mid g_{p,i}(t) \geq v_{j-1}\}$. See Figure 1. It is not hard to show the following lemma.

Lemma 1. *The total (weighted) throughput of a schedule A for the constructed MAX-THP instance equals its total profit when interpreted as a schedule for the original MAX-PFT instance and vice versa.*




Fig. 1. Illustrations of the slicing trick. The left hand side is a request with a general profit function and the right hand side is one with a unimodal profit function.

If each profit function is *unimodal*, meaning that it is non-decreasing up to a point and non-increasing after that point, we observe that the slicing trick should create requests each having only one request interval since the time slots $\{t \mid g_{p,i}(t) > v_{r-1}\}$ are consecutive (see the right hand side of Figure 1). Therefore, we can apply any approximation algorithm that works for the weighted throughput maximization problem with one interval for each request and obtain the same approximation ratio for profit maximization with unimodal profit functions. The best known approximation ratio for weighted throughput maximization is $3/4$ due to Gandhi et al. [17].

Theorem 1. *For arbitrary non-negative unimodal profit functions, there is a $3/4$ -approximation for MAX-PFT.*

However, if the profit function is not unimodal, the resulting MAX-THP instance may contain requests that have multiple request intervals. Next, we show that a simple independent rounding scheme that gives a $(1 - 1/e)$ -approximation for MAX-THP with each request associated with one or more intervals, which implies a $(1 - 1/e)$ -approximation for MAX-PFT.

Let $x_{p,i}, y_p^{(t)}$ be the optimal fractional solution of LP (1). Consider the following simple independent rounding scheme: Consider each time slot t independently and choose exactly one page to broadcast. Page p is chosen with probability $y_p^{(t)}$. Note that this is feasible since $\sum_p y_p^{(t)} \leq 1$. We can easily lower bound the probability that a request is satisfied by the schedule produced by the independent rounding.

Lemma 2. *Using independent rounding, the probability that a request $J_{p,i}$ is satisfied is at least $(1 - 1/e)x_{p,i}$.*

The expected total number of requests captured is thus

$$\sum_{p,i} \Pr[J_{p,i} \text{ is satisfied}] \geq (1 - \frac{1}{e}) \sum_{p,i} w_{p,i} x_{p,i} \geq (1 - \frac{1}{e}) OPT.$$

We thus conclude:

Theorem 2. *For any non-negative profit functions, there is a $(1 - 1/e)$ -approximation for MAX-PFT.*

MAX-PFT via Submodular set function maximization: An alternative algorithm achieving the same ratio can also be obtained by casting MAX-PFT as a special case of the problem of maximizing a monotone submodular set function subject to a matroid constraint. For recent progress on constrained submodular set function maximization, see e.g. [3] and reference therein.

First we give the definition of matroid. Let N be a finite set and \mathcal{I} be a family of subsets of N . The pair (N, \mathcal{I}) is called matroid if \mathcal{I} satisfies the following properties. (1) \mathcal{I} is non-empty, (2) downward closed: if $A \in \mathcal{I}$ and $B \subseteq A$, then $B \in \mathcal{I}$, and (3) independent: if $A, B \in \mathcal{I}$ and $|A| < |B|$, then $A \cup \{x\} \in \mathcal{I}$ for some $x \in B \setminus A$. One special matroid is a partition matroid. In a *partition* matroid, N is partitioned into N_1, N_2, \dots, N_ℓ with associated integers k_1, k_2, \dots, k_ℓ , and $A \in \mathcal{I}$ if and only if $\forall i \mid |A \cap N_i| \leq k_i$.

k_i . Next we give the definition of monotone submodular set function $f : 2^N \rightarrow \mathcal{R}^+$. The function f is called submodular when $f(A + x) - f(A) \leq f(B + x) - f(B)$ whenever $B \subseteq A$ and $x \in N$. By monotonicity, we mean that if $B \subseteq A$ then $f(B) \leq f(A)$, and $f(\emptyset) = 0$. The problem of maximizing the submodular function f under the matroid constraint (N, \mathcal{I}) can be formulated as finding $A = \arg \max_{A' \in \mathcal{I}} f(A')$.

We interpret MAX-PFT as a special case of the above general problem in the following way. Let $N = \mathcal{P} \times \mathcal{T}$, where \mathcal{P} is the set of pages and \mathcal{T} is the set of time slots. Let $N_t = \mathcal{P} \times \{t\}$. Let $A \in \mathcal{I}$ iff $\forall t |A \cap N_t| \leq 1$. Notice that (N, \mathcal{I}) is a partition matroid. The function f is defined as follows: $f = \sum_{p,i} \max_{t:(p,t) \in A} g_{p,i}(t)$. It is not hard to see that f is a monotone submodular set function. It is known that maximizing a monotone submodular function can be approximated with factor $1 - \frac{1}{e}$ under a matroid constraint [3]. Therefore, we can obtain a $1 - \frac{1}{e}$ -approximation for MAX-PFT.

The advantage of the alternative algorithm above is the following. Once the connection to submodular functions and matroid constraints is seen, one can readily obtain similar results for more general settings. For example, it is possible that a client request can be satisfied by sending any one of several similar pages. In this case, as long as, one is able to define an appropriate submodular profit function, one again obtains a $(1 - 1/e)$ -approximation. Moreover, one can also impose additional constraints as long as they satisfy a matroid constraint; multiple matroid constraints can also be handled with some additional loss in the approximation. Finally, one can also obtain concentration bounds in some cases [9] and these can be useful in handling additional constraints. We defer a detailed description of these extensions to a later version of the paper.

A 2-Speed 1-Approximation for Throughput Maximization In this section, our goal is to show a randomized 2-speed 1-approximation for throughput maximization. Here the objective is to satisfy as many requests as possible by their deadlines. Recall that in [6] gave an algorithm to convert a fractional solution that satisfies all requests to a 2-speed integer solution with the same property. To obtain a true 1-approximation, we need to also handle the case where the fractional solution does not satisfy all requests. Our analysis relies on the result of [6]. For completeness, we begin by showing that if there is a feasible fractional solution to LP (1) that satisfies all requests, there is a 2-speed integral scheduling that can also satisfy all requests. Then we show how to extend this to obtain the 2-speed 1-approximation by using dependent rounding.

Let $x_{p,i}, y_p^{(t)}$ be a fractional solution to the LP where all requests are satisfied by their deadlines. We first construct a bipartite graph $G = (U, V, E)$ as follows. One partite set U contains a vertex u_t representing each time slot t . Let $z_p^{(t)} = \sum_{t'=1}^t y_p^{(t')}$ be the cumulative amount of page p transmitted through time t in the fractional solution. Let $I_{p,t} = [z_p^{(t-1)}, z_p^{(t)})$, which represents the fractional occurrences of page p transmitted during time slot t .

The other partite set V contains, for each page p and for $i = 1$ to $\lceil 2z_p^{(T)} \rceil$, a vertex $v_{p,i}$ representing the i th fractional transmission of half of page p . Let $I(v_{p,i}) = [(i-1)/2, i/2)$ and let $W_{p,i}$ be the *window* of consecutive time slots t such that $I_{p,t}$ overlaps $I(v_{p,i})$. For each $t \in W_{p,i}$, we add an edge $(u_t, v_{p,i})$. See Figure 2 for an illustration of the construction.

Now we augment the bipartite graph to get a network flow instance. We add a source \mathbf{s} , edges (\mathbf{s}, u_t) with capacity 1 for $\forall t$, a sink \mathbf{t} and edges $(v_{p,i}, \mathbf{t})$ with capacity $1/2$ for $\forall p, i$. First we can observe that there is \mathbf{s} - \mathbf{t} flow f with flow value $\sum_p z_p^{(T)}$. In fact, just by letting $f(u_t, v_{p,i})$ be the measure (i.e., length) of $I_{p,t} \cap I(v_{p,i})$ and setting $f(\mathbf{s}, u_t) \forall t$ and $f(v_{p,i}, \mathbf{t}) \forall i$ accordingly, f is such a flow. For each page p , we delete the last vertex $v_{p, \lceil 2z_p^{(T)} \rceil}$ if $f(v_{p, \lceil 2z_p^{(T)} \rceil}, \mathbf{t}) < 1/2$. After this, we can see that f saturates all edges $(v_{p,i}, \mathbf{t}) \forall p, i$.

Next, we double the capacities of all edges and find a maximum \mathbf{s} - \mathbf{t} integral flow f' . This is possible since all capacities are now integral. The obtained integral flow can be interpreted as a 2-speed schedule: If there is a unit of flow going from u_t to $v_{p,i}$, the server will broadcast p at time t . Since the capacity of (\mathbf{s}, u_t) is 2, at most 2 pages are broadcast in one time slot. Note that all edges $(v_{p,i}, \mathbf{t})$ are saturated. This in turn means that for each window $W_{p,i}$, the server broadcasts p at least once in some time slot $t \in W_{p,i}$. For each request $J_{p,i}$, we know that $\sum_{t \in \mathcal{T}_{p,i}} y_p^{(t)} \geq 1$. Therefore, some window $W_{p,j}$ is fully contained in $\mathcal{T}_{p,i} = \{a_{p,i}, \dots, d_{p,i}\}$ due to the construction of the windows. Hence, all requests are satisfied by the 2-speed schedule.

Now, we generalize the above idea to get a true 2-speed 1-approximation, that is a schedule such that the server broadcasts at most 2 pages in one time slot and satisfies at least the number of requests that can be satisfied by the optimum 1-speed schedule. The idea is very simple, instead of scaling the capacities, we just take the bipartite graph G and the flow f , and do dependent rounding on G with values $2 \times f$. We notice that all f values defined on the edges of G are at most $1/2$. Therefore, $2 \times f$ are valid probabilities. Consider a request $J_{p,i}$ which is not completely satisfied, i.e., $\sum_{t \in \mathcal{T}_{p,i}} y_p^{(t)} < 1$.

In this case, $x_{p,i} = \sum_{t \in \mathcal{T}_{p,i}} y_p^{(t)}$. It is easy to see that $\mathcal{T}_{p,i}$ is fully contained in two windows $W_{p,j}, W_{p,j+1}$ for some j . Let $y = \sum_{t \in \mathcal{T}_{p,i}} f(u_t, v_{p,j})$. By **(P2)** of the dependent rounding scheme, we know at most 1 edge incident on a window can be chosen. Therefore, for fixed p, j , the events that $(u_t, v_{p,j})$ is chosen are disjoint. Then, by **(P1)**,

$$\Pr((u_t, v_{p,j}) \text{ is chosen for some } t \in \mathcal{T}_{p,i}) = 2 \sum_{t \in \mathcal{T}_{p,i}} y_p^{(t)} = 2y.$$

Similarly, we can show that $\Pr((u_t, v_{p,j+1}) \text{ is chosen for some } t \in \mathcal{T}_{p,i}) = 2(x - y)$. Therefore, $\Pr(J_{p,i} \text{ is satisfied}) \geq \max(2y, 2(x - y)) \geq x_{p,i}$. If $J_{p,i}$ is completely satisfied, we can use the previous argument, that is $\mathcal{T}_{p,i}$ fully contains some window $W_{p,j}$ and some edge incident to $v_{p,i}$ must be chosen. Again, we have $\Pr(J_{p,i} \text{ is satisfied}) = 1 = x_{p,i}$. Since we have shown that each request is satisfied with a probability no smaller than the probability that the request is satisfied in the fractional optimal solution, we obtain the following theorem.

Theorem 3. *There is a polynomial time 2-speed 1-approximation for MAX-THP.*

Throughput Maximization with a Relaxed Time Window In this section, we assume that each request is fractionally fully satisfied by the optimal solution of LP(1), i.e., $x_{p,i} = 1 \forall p, i$. Suppose a request $J_{p,i}$ arrives at time $a_{p,i}$ with deadline $d_{p,i}$ (associated with the window $[a_{p,i}, d_{p,i}]$), then we construct an integral schedule such that this request is satisfied within the window $[a_{p,i} - l_{p,i}, d_{p,i} + l_{p,i}]$ where $l_{p,i} = d_{p,i} - a_{p,i} + 1$




Fig. 2. The construction of the bipartite graph $G(U, V, E)$. E.g. $I_{p,4} = [1.2, 1.75)$, $I(v_{p,4}) = [1.5, 2)$ and $I((u_4, v_{p,4})) = I_{p,4} \cap I(v_{p,4}) = [1.5, 1.75]$.

is the length of the window $\mathcal{T}_{p,i}$. By left shifting the window or right shifting the window by its length, we can satisfy the request. A shifting, or expanding of the window is necessary and we refer to this as a relaxed schedule since it satisfies all the requests in a relaxed manner, and the client request is satisfied at a time approximately close to the desired window of time. For a given instance, we consider a fractional solution which, by assumption, satisfies all requests before their deadlines.

Starting from the instance I that has a fractional solution in which every request is satisfied, we will create an instance \mathcal{I} which is a subset of the requests such that finding an integral solution for \mathcal{I} will also immediately lend a *relaxed integral solution* to the instance \mathcal{I} . We focus on a single page p . Order all the client requests for page p in order of non-decreasing window length. Initially \mathcal{I} is the empty set of requests. We try to insert each request (in non-decreasing window length order) into set \mathcal{I} , and as long as it does not overlap with a request already inserted into \mathcal{I} , we insert it. This will give us a collection of non-overlapping requests for page p . We do this filtering for every possible page. This gives us a fractional solution in which all requests for the same page are non-overlapping and completely satisfied. Using flow based methods⁶ it is easy to convert this to an integral solution that satisfies all the requests. Client requests in \mathcal{I} are clearly satisfied (integrally) within their intervals. Each client request $J_{p,i}$ that was not chosen in \mathcal{I} overlapped with a chosen request with a smaller window size. Thus it is also satisfied in the integral solution within time $[a_{p,i} - l_{p,i}, d_{p,i} + l_{p,i}]$, i.e., satisfied within the relaxed deadlines. We thus conclude:

Theorem 4. *Suppose there is a fractional schedule that satisfies all requests. We can convert the fractional solution to an integral one in polynomial time such that each request $J_{p,i}$ can be satisfied in the relaxed window $[a_{p,i} - l_{p,i}, d_{p,i} + l_{p,i}]$ where $l_{p,i} = d_{p,i} - a_{p,i} + 1$ is the length of the window $\mathcal{T}_{p,i}$.*

3.2 Online Algorithms

In this section we revisit the problem MAX-PFT, but now in the *online* setting, in which a request is not known to the server until it arrives. As previously discussed in Section 3, maximizing the total profit can be interpreted as maximizing a monotone submodular function subject on a matroid. It is known that a simple greedy algorithm gives

⁶ This involves the same technique as used for converting a fractional matching in a bipartite graph to an integral matching [16].

2-approximation for such a problem [24]. Further, the greedy algorithm can be interpreted as an online algorithm in this setting. Thus we can easily obtain a 2-competitive algorithm for MAX-PFT. For the more restricted setting MAX-THP, [21] gave a 2-approximation. Here we show that the greedy algorithm’s performance improves in the resource augmentation model when the algorithm is given a speed larger than 1. There is no natural way to interpret resource augmentation in the general framework of submodular set function maximization subject to a matroid constraint. We therefore resort to a direct analysis.

We will be considering a resource augmentation analysis [19]. In this analysis, the online algorithm is given s -speed and compared to a 1 speed optimal solution. For some objective function, we say that an algorithm is s -speed c -competitive if the algorithm’s objective is within a c factor of the optimal solution’s objective. We describe our greedy algorithm *Maximum Additional Profits First* (for short, MAPF) which is given an integer speed $s \geq 1$. As implied in its name, the algorithm MAPF broadcasts s pages which give the maximum additional profits by broadcasting.

Algorithm: MAPF

- At any time t , broadcast s pages which give the maximum additional profits.

For this algorithm we show the following theorems. Proofs are deferred to a full version of this paper.

Theorem 5. MAPF is s -speed $(1 + 1/s)$ -competitive online algorithm for MAX-PFT and MAPF is s -speed $(1 + 1/s)$ -competitive algorithm for MAX-THP.

Theorem 6. For any $\epsilon > 0$ and speed $s \geq 1$, MAPF is not s -speed $(1 + 1/s - \epsilon)$ -competitive for MAX-THP.

For $s = 1$, for any $\epsilon > 0$, there is a lower bound of $(2 - \epsilon)$ on the competitive ratio of any online algorithm, even if it is randomized [21].

References

1. N. Bansal, D. Coppersmith, and M. Sviridenko. Improved approximation algorithms for broadcast scheduling. *SODA*, pp. 344–353, 2006.
2. Y. Bartal and S. Muthukrishnan. Minimizing maximum response time inscheduling broadcasts. *SODA*, pp. 558–559, 2000.
3. G. Calinescu, C. Chekuri, M. Pal, and J. Vondrak. Maximizing a monotone submodular set function subject to a matroid constraint *SIAM J. on Computing*, to appear.
4. W. Chan, T. Lam, H. Ting, and P. Wong. New Results on On-Demand Broadcasting with Deadline via Job Scheduling with Cancellation. *COCOON*, pp. 210–218, 2004.
5. D. Carney, S. Lee and S. Zdonik. Scalable Application-Aware Data Freshening. *ICDE*, pp. 481–492 (2003).
6. J. Chang, T. Erlebach, R. Gailis and S. Khuller. Broadcast Scheduling: Algorithms and Complexity. *SODA*, pp. 473–482, 2008.
7. M. Charikar, S. Khuller. A robust maximum completion time measure for scheduling. *SODA*, pp. 324–333, 2006.
8. C. Chekuri, S. Im and B. Moseley. Minimizing Maximum Response Time and Delay Factor in Broadcast Scheduling. *ESA*, pp. 444–455, 2009.

9. C. Chekuri, J. Vondrak and R. Zenklusen. Dependent Randomized Rounding via Exchange Properties of Combinatorial Structures. *FOCS*, 2010.
10. M. Chrobak, C. Dürr, W. Jawor, L. Kowalik and Maciej Kurowski. A Note on Scheduling Equal-Length Jobs to Maximize Throughput. *J. of Scheduling*, Vol. 9(1):71–73, (2006).
11. P. Deolasee, A. Katkar, P. Panchbudhe, K. Ramamritham and P. Shenoy. Adaptive Push-Pull: Disseminating Dynamic Web Data. *WWW*, (2001).
12. J. Eckstein, A. Gal and S. Reiner. Monitoring an Information Source under a Politeness Constraint. *INFORMS Journal on Computing*, Vol. 20(1):3–20, (2007).
13. J. Edmonds and K. Pruhs. Multicast pull scheduling: when fairness is fine. *SODA*, pp. 421–430, 2002.
14. J. Edmonds and K. Pruhs. A maiden analysis of longest wait first. *SODA*, pp. 811–820, 2004.
15. A. Gal and J. Eckstein. Managing Periodically Updated Data in Relational Databases: A Stochastic Modeling Approach. *JACM*, Vol. 48(6):1141–1183, (2001).
16. R. Gandhi, S. Khuller, Y. Kim, and Y.C. Wan. Algorithms for minimizing response time in broadcast scheduling. *IPCO*, pp. 415–424.
17. R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. Dependent rounding and its applications to approximation algorithms. *JACM*, Vol 53(3):324–360, (2006). Preliminary version in *FOCS*, 2002.
18. S. Im and B. Moseley. An Online Scalable Algorithm for Average Flow Time in Broadcast Scheduling. *SODA*, 2010.
19. B. Kalyanasundaram and K. Pruhs. Speed is as powerful as clairvoyance. *J. ACM*, Vol 47(4):617–643, (2000).
20. B. Kalyanasundaram, K. Pruhs, and M. Velauthapillai. Scheduling broadcasts in wireless networks. *ESA*, pp. 290–301, 2000.
21. J. Kim and K. Chwa. Scheduling broadcasts with deadlines. *Theor. Comput. Sci.*, Vol 325:479–488, (2004).
22. R. Motwani and P. Raghavan. Randomized Algorithms. *ACM Comput. Surveys*, Vol. 28(1):33–37 (1996).
23. R. Motwani and P. Raghavan. Randomized Algorithms. Cambridge University Press, 1995.
24. G. L. Nemhauser, L. A. Wolsey and M. L. Fisher. An analysis of approximations for maximizing submodular set functions - I. *Mathematical Programming*, Vol 14(1):265–294, (1978).
25. S. Pandey, K. Dhamdhere and C. Olston. WIC: A General-Purpose Algorithm for Monitoring Web Information Sources. *VLDB*, pp. 360–371, (2004).
26. H. Roitman, A. Gal and L. Raschid. Satisfying Complex Data Needs using Pull-Based Online Monitoring of Volatile Data Sources. *ICDE*, 2008.
27. H. Roitman. Profile Based Online Data Delivery - Model and Algorithms. *Ph.D. Thesis* (2008).
28. H. Roitman, A. Gal, L. Raschid. On Trade-offs in Event Delivery Systems. The 4th ACM International Conference on Distributed Event-Based Systems (DEBS), 2010
29. F. Zheng, S. Fung, W. Chan, F. Chin, C. Poon and P. Wong. Improved On-Line Broadcast Scheduling with Deadlines. *COCOON*, 320–329, 2006.