

Pull-Based Data Broadcast with Dependencies: Be Fair to Users, not to Items

Julien Robert*

Nicolas Schabanel*,†

Abstract

Broadcasting is known to be an efficient means of disseminating data in wireless communication environments (such as Satellite, mobile phone networks,...). It has been recently observed that the average service time of broadcast systems can be considerably improved by taking into consideration existing correlations between requests. We study a pull-based data broadcast system where users request possibly overlapping sets of items; a request is served when all its requested items are downloaded. We aim at minimizing the average user perceived latency, *i.e.* the average flow time of the requests. We first show that any algorithm that ignores the dependencies can yield arbitrary bad performances with respect to the optimum even if it is given arbitrary extra resources. We then design a $(4 + \epsilon)$ -speed $O(1 + 1/\epsilon^2)$ -competitive algorithm for this setting that consists in 1) splitting evenly the bandwidth among each requested set and in 2) broadcasting arbitrarily the items still missing in each set into the bandwidth the set has received. Our algorithm presents several interesting features: it is simple to implement, non-clairvoyant, fair to users so that no user may starve for a long period of time, and guarantees good performances in presence of correlations between user requests (without any change in the broadcast protocol). We also present a $(4 + \epsilon)$ -speed $O(1 + 1/\epsilon^3)$ -competitive algorithm which broadcasts at most one item at any given time and preempts each item broadcast at most once on average. As a side result of our analysis, we design a competitive algorithm for a particular setting of non-clairvoyant job scheduling with dependencies, which might be of independent interest.¹

Keywords: Multicast scheduling, Pull-based broadcast, Correlation-based, Non-clairvoyant schedul-

ing, Resource augmentation.


1 Introduction

Motivations. Broadcasting is known to be an efficient means of disseminating data in wireless communication environments (such as Satellite, mobile phone networks,...). It has been recently observed in [13, 12, 5] that the average service time of broadcast systems can be considerably improved by taking into consideration existing correlations between requests. Most of the theoretical research on data broadcasting was conducted until very recently under the assumption that user requests are for a single item at a time and are independent of each other. However, users usually request several items at a time which are, to a large extent, correlated. A typical example is a web server: users request web pages that are composed of a lot of shared components such as logos, style sheets, title bar, news headers,..., and all these components have to be downloaded together when any individual page is requested. Note that some of these components, *e.g.* news header, may constantly vary over time (size and/or content).

Pull-based data broadcast with dependencies. We study a pull-based data broadcast system where users request possibly overlapping sets of items. We aim at minimizing the average user perceived latency, *i.e.* the *average flow time* of the requests, where the flow time of a request is defined as the time elapsed between its arrival and the end of the download of the last requested item. We assume that user cannot start downloading an item in the middle of its broadcast. When the broadcast of an item starts, all the outstanding requests asking for this item can start downloading it. Several items may be downloaded simultaneously. We consider the *online* setting where the scheduler is *non-clairvoyant* and discovers each request at the time of its arrival; furthermore, the scheduler does not even know the lengths of the requested items and is aware of the completion of a broadcast only at the time of its completion. Items are however labeled with a unique ID to allow their retrieval. Note that this are the typical requirements of a real life systems where items may vary over time.

*École normale supérieure de Lyon, Laboratoire de l'informatique du parallélisme, UMR CNRS ENS-LYON INRIA UCBL n°5668, 46 allée d'Italie, 69364 Lyon Cedex 07, France, <http://perso.ens-lyon.fr/julien.robert>.

†CNRS, Centro de Modelamiento Matemático, UMI CNRS Universidad de Chile n°2807, Blanco Encalada 2120 Piso 7, Santiago de Chile, Chile, <http://cmm.uchile.cl/~schabanel>.

¹This work was supported by the CNRS Grant .

Background. It is well known that preemption is required in such systems in order to achieve reasonable performances. Furthermore, [8] proved that even without dependencies, no algorithm can guarantee a flow time less than $\Omega(\sqrt{n})$ times the optimal. The traditional approach in online algorithms consists then in penalizing the optimum by increasing the bandwidth given to the algorithm so that its performances can be compared to the optimum. This technique is known as *resource augmentation* and provides interesting insights on the relative performances of different algorithms that could not be compared directly to the optimum cost. In our case, we give to our algorithm a bandwidth $s > 1$ and show that it achieves a flow time less than a constant times the optimum cost with a bandwidth 1. Formally, an algorithm is *s-speed c-competitive* if when given a bandwidth s , its flow time is at most at a factor c of the optimum flow time with bandwidth 1.

To our knowledge the only positive results [8, 9] in the online setting assume that the requests are independent and ask for one single item. The authors show that without dependencies the algorithms **EQUI** and **LWF** are competitive. **EQUI** which splits evenly the bandwidth among the alive requested items, is $(4 + \epsilon)$ -speed $(2 + 8/\epsilon)$ -competitive, and **LWF**, which broadcasts the item where the aggregate waiting times of the outstanding requests for that item is maximized, is 6-speed $O(1)$ -competitive (where the bound proved on the competitive ratio is $O(1) = 6,000,000$). In the *offline* setting, where the requests and their arrival times are known at time $t = 0$, the problem is already NP-hard but better bounds can be obtained using linear programming [14, 10, 11, 1, 2]; the latest result, [2] to our knowledge, is a $O(\log^2(T + n)/\log \log(T + n))$ -approximation where n is the number of requests and T the arrival time of the last request. To our knowledge, our results are the first provably efficient algorithms to deal with dependencies in the online setting.

Concerning the push-based variant of the problem, where the requests arrival times follow some Poisson process and the requested sets are identically distributed according to a fixed distribution, constant factor approximations exist in presence of dependencies [4, 3, 6]. The latest result, [6], obtains a 4-approximation if the requested sets are drawn according to an arbitrary fixed distribution over a finite number of subsets of items.

Our contribution. We first show that the performances of any algorithm that ignores the dependencies can be arbitrarily far from the optimal cost even if it is given *arbitrary* extra resources. We then design a $(4 + \epsilon)$ -speed $O(1 + 1/\epsilon^2)$ -competitive algorithm **B-EQUISET** for the non-clairvoyant data broadcast problem with dependencies. **B-EQUISET** consists in 1) splitting evenly

the bandwidth among each requested set and in 2) broadcasting arbitrarily the items still missing in each set into the bandwidth the set has received. The spirit of the algorithm is that *one should favor the users over the items* in the sense that it splits the bandwidth evenly among the outstanding requested sets and arbitrarily among the outstanding items within each requested set. Our algorithm presents several interesting features: it is simple to implement, non-clairvoyant, fair to users so that no user may starve for a long period of time, and improves performances in presence of correlations between user requests (without any change in the broadcast protocol). Precisely, we prove that:

THEOREM 1.1. (MAIN RESULT) *For all $\delta > 0$ and $\epsilon > 0$, **B-EQUISET** is a $(1 + \delta)(4 + \epsilon)$ -speed $(2 + 8/\epsilon)(1 + 1/\delta)$ -competitive algorithm for the online data broadcast problem with dependencies.*

One could object that **B-EQUISET** is unrealistic since it can split the bandwidth arbitrarily. But using the same technique as in [8], it is easy to modify **B-EQUISET** to obtain an other competitive algorithm **B-EQUISET-EDF** (described at the end of section 5) which, with a slight increase of bandwidth, ensures that at most one item is broadcast at any given time and that each broadcast is preempted at most once on average.

THEOREM 1.2. (BOUNDED PREEMPTION) *For all $\delta > 0$ and $\epsilon > 0$, **B-EQUISET-EDF** is a $(1 + \delta)^2(4 + \epsilon)$ -speed $(2 + 8/\epsilon)(1 + 1/\delta)^2$ -competitive algorithm for the online data broadcast problem with dependencies, where each broadcast is preempted at most once on average.*

Our analysis takes its inspiration in the methods developed in [8]. In order to extend their analysis to our algorithm, we have also designed a new competitive algorithm **EQUI \circ A** for a particular setting of non-clairvoyant job scheduling with dependencies which might be of independent interest (Theorem 4.1).

The next section gives a formal description of the problem and shows that it is required to take dependencies into account to obtain a competitive algorithm. Section 3 exposes the algorithm **B-EQUISET** and introduces useful notations. Section 4 designs a competitive algorithm **EQUI \circ A** for a variant of job scheduling with dependencies that is used in Section 5 to analyze the competitiveness of our algorithm **B-EQUISET**.

2 Definitions and notations

The problem. The input consists of:

- A set \mathcal{J} of n items I_1, \dots, I_n of length ℓ_1, \dots, ℓ_n , respectively

- A set \mathcal{S} of q requests for q non-empty sets of items $S_1, \dots, S_q \subseteq \mathcal{J}$, with arrival times a_1, \dots, a_q , respectively.

Schedule. A s -speed schedule is an allocation of a bandwidth of size s to the items of \mathcal{J} over the time. Formally, it is described by a function $r : \mathcal{J} \times [0, \infty) \rightarrow [0, s]$ such that for all times t , $\sum_{I \in \mathcal{J}} r(I, t) \leq s$; $r(I, t)$ represents the *rate* of the broadcast of I at time t , *i.e.*, the amount of bandwidth allotted to item I at time t . An item I_i is broadcast between t and t' if its broadcast starts at time t and if the total bandwidth allotted to I_i between t and t' sums up to ℓ_i , *i.e.*, if $\int_t^{t'} r(I_i, t) dt = \ell_i$. We denote by $c(I_i, k)$ the date of the completion of the k th broadcast of item I_i . Formally, it is the first date such that $\int_0^{c(I_i, k)} r(I_i, t) dt = k \ell_i$ (note that $c(I_i, 0) = 0$). We denote by $b(I_i, k)$ the date of the beginning of the k th broadcast of item I_i , *i.e.* $b(I_i, k) = \inf\{t \geq c(I_i, k-1) : r(I_i, t) > 0\}$.²

Cost. For all times t , let $B(I_i, t)$ be the time of the beginning of the first broadcast of item I_i after t , *i.e.* $B(I_i, t) = \min\{b(I_i, k) : b(I_i, k) \geq t\}$. For all times t , $C(I_i, t)$ denotes the time of the end of the first broadcast of item I_i starting after t , *i.e.* $C(I_i, t) = \min\{c(I_i, k) : b(I_i, k) \geq t\}$. The *completion time* c_j of request S_j is the first time such that every item in S_j has been broadcast (or downloaded) after its arrival time a_j , *i.e.*, $c_j = \max_{I_i \in S_j} C(I_i, a_j)$. We aim at minimizing the *average completion time* defined as $\frac{1}{q} \sum_{S_j \in \mathcal{S}} (c_j - a_j)$, or equivalently the *flow time* defined as the sum of the waiting times, *i.e.* $\text{B-FlowTime} = \sum_{S_j \in \mathcal{S}} (c_j - a_j)$. We denote by $\text{BOPT}_s(\mathcal{S})$ the flow time of an optimal s -speed schedule for a given instance \mathcal{S} .

s -Speed c -Competitive Algorithms. We consider the online setting of the problem, in which the scheduler gets informed of the existence of each request S_j at time a_j and not before. The scheduler is not even aware of the lengths $(\ell_i)_{I_i \in S_j}$ of the requested items in each set nor of the total number n of available items. It is well known (*e.g.*, see [8]) that in this setting, it is impossible to approximate within a factor $o(\sqrt{n})$ the optimum flow time for a given bandwidth s even if all items have unit length (independently of any conjecture such as $P = NP$). The traditional approach in online algorithms consists then in penalizing the optimum by

²Remark that this formalization prevents from broadcasting the same item twice at a given time or from aborting the current broadcast of an item. The first point is not restrictive since if two broadcasts of the same item overlap, one reduces the service time by using the beginning of the bandwidth allotted to the second broadcast to complete earlier the first, and then the end of the first to complete the second on time. The second point is at our strict disadvantage since it does not penalize an optimal schedule that would never start a broadcast to abort it later on.

increasing the bandwidth given to the algorithm so that its performances can be compared to the optimum. This technique is known as *resource augmentation* and provides interesting insights on the relative performances of different algorithms that could not be compared directly to the optimum cost. In our case, we give to our algorithm a bandwidth $s > 1$ and show that it achieves a flow time less than a constant times the optimum cost with a bandwidth 1. Formally, an algorithm is s -speed c -competitive if when given s times as many resources as the adversary, its cost is no more than c times the optimum cost. In our case the resource is the bandwidth, and we compare the cost A_s of a scheduler A with a bandwidth s , to the cost BOPT_1 of an optimal schedule on a unit bandwidth. (We denote by A_s the cost of an algorithm A when given a bandwidth s .)

We show below that any algorithm ignoring existing dependencies (*i.e.*, an algorithm that chooses to ignore the actual set of items asked by each request) can lead to arbitrarily bad solutions. A natural example of such algorithm consists in creating a request per item requested in each set and running previously known algorithms (such as [8, 9]). The following fact shows that this natural algorithm which designs a schedule only based on the number of requests waiting for each item, cannot be competitive.

FACT 2.1. (DEPENDENCIES CANNOT BE IGNORED)

No algorithm A that ignores dependencies is s -speed c -competitive for any $c < \frac{2}{3s} \sqrt{n}$ if A is deterministic, and for any $c < \frac{1}{6s} \sqrt{n}$ if A is randomized.

Proof. Consider first a deterministic algorithm A which is given a bandwidth s and consider the instance where n different items are requested at time $t = 0$. Since A ignores the dependencies, we set them after the execution of the algorithm A : one request asks for the $n - \sqrt{n}$ items that have been served the most by A at time $t = (n - \sqrt{n})/s$, and \sqrt{n} requests ask for each of the remaining \sqrt{n} items. Then, algorithm A serves each request only after time $t = (n - \sqrt{n})/s$ and its flow time is at least $(\sqrt{n} + 1)(n - \sqrt{n})/s \sim n\sqrt{n}/s$. The optimal solution with bandwidth only 1 first broadcasts the items corresponding to the \sqrt{n} unit length requests and then broadcasts the $n - \sqrt{n}$ remaining items; the optimal flow time is then $(n + \sum_{k=1}^{\sqrt{n}} k) \sim \frac{3}{2}n$. This shows a gap of $\frac{2}{3s} \sqrt{n}$ between the optimal cost with bandwidth 1 and every deterministic algorithm with bandwidth $s = O(\sqrt{n})$, which ignores the dependencies.

We use Yao's principle (see [16, 15]) to extend the result to randomized algorithms. We consider the following probabilistic distribution of requests set over n items: $1 + \sqrt{n}$ requests arrive at time $t = 0$; one request asks for an uniform random subset S_0 of size $n - \sqrt{n}$ of

the n items; and each of the requests $S_1, \dots, S_{\sqrt{n}}$ asks for one random distinct item among the \sqrt{n} remaining items. Consider again any deterministic algorithm A with bandwidth s . Since A is deterministic and ignores the dependencies, the schedule designed by A is independent of the random instance. At time $t = n/(2s)$, the broadcast of at least $n/2$ items is not completed. Thus, the probability that request S_j , for $j \geq 1$, asks for one of these items is at least $1/2$. Then, the expected number of unsatisfied request at time $t = n/(2s)$ is at least $\sqrt{n}/2$. We conclude that the expected flow time for any deterministic algorithm with bandwidth s under this distribution of request is at least $n\sqrt{n}/(4s)$. According to Yao's principle, the worst expected flow time of any randomized algorithm over the collection of all the considered instances is at least $n\sqrt{n}/(4s)$. But $\text{BOPT}_1 \sim \frac{3}{2}n$, which concludes that no randomized algorithm is s -speed c -competitive, for all s and $c < \sqrt{n}/(6s)$. \square

3 The Algorithm **B-EQUISET**

Definitions. A request S_j for a subset of items is said to be *alive* at time t if $t \geq a_j$ and if the download of at least one item $I_i \in S_j$ is not yet completed at time t , *i.e.*, $t < C(I_i, a_j)$. We say that an item $I_i \in S_j$ whose download is not yet completed (*i.e.*, such that $a_j \leq t < C(I_i, a_j)$) is *alive for* S_j at time t .

The **B-EQUISET Algorithm.** Consider that we are given a bandwidth s . Let $R(t)$ be the set of alive requests at time t during the execution of the algorithm. For all t , **B-EQUISET** allocates to each alive request the same amount of bandwidth, $s/|R(t)|$; then, for each alive request S_j , it splits *arbitrarily* the $s/|R(t)|$ bandwidth allotted to S_j among its alive items. Precisely, it allocates to each item I_i alive for S_j at time t , an *arbitrary* amount of bandwidth, $r_{j,i}(t) \geq 0$, such that

$$\sum_{I_i \text{ alive for } S_j} r_{j,i}(t) = s/|R(t)|.$$

B-EQUISET then broadcasts at time t each item I_i at a rate

$$r_i(t) = \sum_{S_j \in R(t): I_i \text{ is alive for } S_j \text{ at time } t} r_{j,i}(t).$$

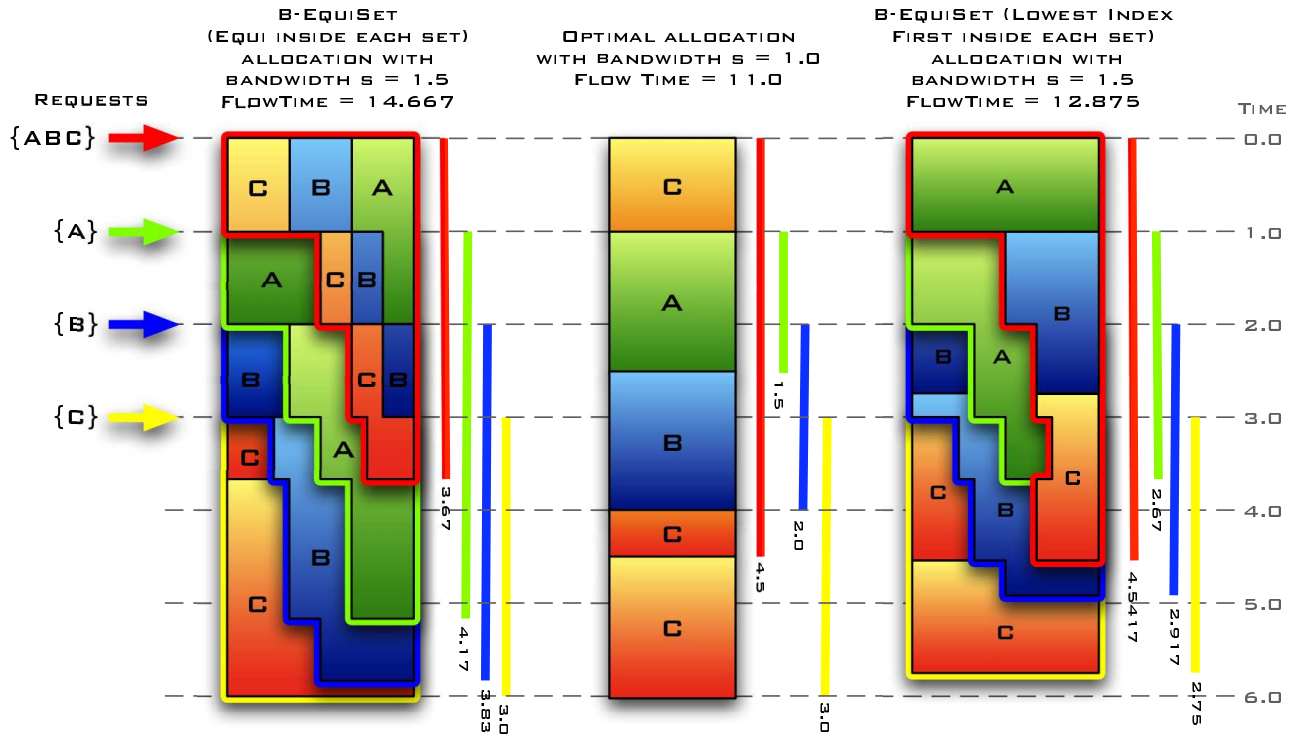
Figure 1 illustrates two executions of the algorithm: in the first (to the left) **B-EQUISET** chooses for each alive request S_j , to divide up the bandwidth allotted to S_j equally among every S_j 's alive items; in the second (to the right), **B-EQUISET** chooses for each alive request S_j , to broadcast in the bandwidth allotted to S_j the lowest indexed alive item in S_j .

Note that bandwidth adjustments for each item are necessary only when new requests arrive or when the broadcast of some item completes.

As in [8], we deduce the performances of our broadcast algorithm **B-EQUISET** from the analysis of the performances of an other algorithm, **EQUI** \circ **A**, for a variant of the non-clairvoyant scheduling problem studied in [7] which includes dependencies. Section 4 presents this later problem and analyzes the competitiveness of algorithm **EQUI** \circ **A**. Then, Section 5 deduces the competitiveness of **B-EQUISET** by simulating **EQUI** \circ **A** on a particular instance of non-clairvoyant scheduling built on the execution of **B-EQUISET**.

4 Non-Clairvoyant Seq-Par Batch Scheduling

For the sake of completeness we first sum up the results in [7], reader may skip this paragraph in a first reading. Edmonds's non-clairvoyant scheduling problem consists in designing an online algorithm that schedules jobs on p processors without any knowledge of the progress of each job before its completion. An instance of non-clairvoyant job scheduling problem consists in a collection of jobs (J_k) with arrival times (a_k); each job J_k goes through a series of phases $J_k^1, \dots, J_k^{m_k}$; the amount of work in each phase J_k^l is w_k^l ; at time t , the algorithm allocates to each uncompleted job J_k an amount ρ_k^t of processors (the (ρ_k^t) s are arbitrary non-negative real numbers, such that at any time: $\sum_k \rho_k^t \leq p$); each phase J_k^l progresses at a rate given by a speed-up function $\Gamma_k^l(\rho_k)$ of the amount ρ_k of processors allotted to J_k during phase J_k^l , that is to say that the amount of work accomplished between t and $t + dt$ during phase J_k^l is $\Gamma_k^l(\rho_k^t)dt$; let t_k^l denote the completion time of the l -th phase of J_k , *i.e.* t_k^l is the first time t' such that $\int_{t_k^{l-1}}^{t'} \Gamma_k^l(\rho_k^t) dt = w_k^l$ (with $t_k^0 = a_k$). The overall goal is to minimize the *flow time* of the jobs, that is to say the sum of the processing time of each job, *i.e.* $\text{J-FlowTime} = \sum_k (t_k^{m_k} - a_k)$. We denote by $\text{JOPT}_s(\mathcal{J})$ the flow time of an optimal s -speed schedule for \mathcal{J} . The algorithm is *non-clairvoyant* in the sense that it does not know anything about the progress of each job and is only informed that a job is completed at the time of its completion. In particular, it is not aware of the different phases that the job goes through (neither of the amount of work nor of the speed-up function). One of the striking results of [7] is that in spite of this total lack of knowledge, the algorithm **EQUI** that allocates an equal amount of processors to each uncompleted job is $(2 + \epsilon)$ -speed $(2 + 4/\epsilon)$ -competitive when the speed up functions are arbitrary non-decreasing sub-linear functions (*i.e.*, such that for all $\rho < \rho'$, $\Gamma_k^l(\rho)/\rho \geq \Gamma_k^l(\rho')/\rho'$, for all k, l).



THE INSTANCE consists of three items A, B, C of length 1.5 and four requests $S_1 = \{A, B, C\}$ (in red), $S_2 = \{A\}$ (in green), $S_3 = \{B\}$ (in blue), and $S_4 = \{C\}$ (in yellow) with arrival times $a_1 = 0$, $a_2 = 1$, $a_3 = 2$, and $a_4 = 3$.

THREE SCHEDULES are presented: an optimal schedule with unit bandwidth (at the center) and two executions of **B-EQUISET** with bandwidth $s = 1.5$. To the left, **B-EQUISET** evenly splits the bandwidth of each alive requested set among its alive items. To the right, **B-EQUISET** gives all the bandwidth allotted to each set to its lowest indexed alive item.

TIME FLIES DOWNWARDS. Four lines to the right of each schedule represent each request's lifetime; the bandwidth allotted to each request is outlined in its respective color. The fill color of each item gets darker as its broadcast progresses.

CONSIDER, FOR INSTANCE, the schedule to the left. **B-EQUISET** first allots all the bandwidth to S_1 and splits it evenly among its items A, B and C (items A, B , and C get darker and darker as their broadcasts progress). At time 1, S_2 arrives and **B-EQUISET** splits the bandwidth evenly between S_1 and S_2 , thus item A is broadcast at a rate: $1.5 \times (\frac{1}{2} + \frac{1}{2} \times \frac{1}{3}) = 1$ and its broadcast completes at time 2. At time 2, S_3 arrives, and **B-EQUISET** splits the bandwidth evenly between S_1, S_2 and S_3 ; S_1 has completed its download of A , thus **B-EQUISET** splits the bandwidth allotted to S_1 among B and C only; S_2 was too late to download A , so it starts a new broadcast of A . S_1, S_2, S_3 , and S_4 are finally served at time $3 + \frac{2}{3}$, $5 + \frac{1}{6}$, $5 + \frac{5}{6}$ and 6, for a total flow time $\mathbf{B-EQUISET}_{1.5}(S) = 14 + \frac{2}{3}$ whereas $\mathbf{BOPT}_1 = 11$. Note that the lowest-index-first execution of **B-EQUISET** (to the right) does better for this instance: $\mathbf{B-EQUISET}'_{1.5}(S) = 12 + \frac{7}{8}$.

Figure 1: Two 1.5-speed executions of **B-EQUISET** algorithm.

Two particular kinds of phases are of interest for our purposes: sequential and parallel. During a *sequential* phase, $\Gamma(\rho) = 1$, that is to say that the job progresses at a unit rate whatever amount of processing power it receives (even if it receives no processor at all, *i.e.* even if $\rho = 0$)! During a *parallel* phase, the job progresses proportionally to the processing power it receives, *i.e.* $\Gamma(\rho) = \rho$. Remark that these two kinds of speed-up functions match the requirement of Edmond’s theorem and thus **EQUI** is $(2 + \epsilon)$ -speed $(2 + 4/\epsilon)$ -competitive on instances consisting of a collection of jobs composed of sequential and parallel phases.

As in [8], we reduce the analysis of our broadcast algorithm **B-EQUISET** to the analysis of a non-clairvoyant scheduling algorithm. For that purpose, we need to introduce dependencies between the jobs in Edmond’s framework. We consider the following variant of the non-clairvoyant scheduling problem.

Non-Clairvoyant Seq-Par Batches Scheduling. An instance of this variant consists in a collection $\mathcal{B} = \{B_1, \dots, B_q\}$ of batches $B_j = \{J_{j,1}, \dots, J_{j,u_j}\}$ of jobs with arrival times a_1, \dots, a_q , where each job $J_{j,i}$ is composed of two phases: a sequential phase of work $w_{j,i}^s \geq 0$ followed by a parallel phase of work $w_{j,i}^p \geq 0$. (Note that this problem is different from the classical batch scheduling problem in which *only one* batch has to be treated.) The scheduler is non-clairvoyant and discovers each batch of jobs at the time of its arrival and is in particular *not aware of the amounts of work of each job in each batch*. The scheduler allocates to each job $J_{j,i}$, arrived and uncompleted at time t , a certain amount $\rho_{j,i}^t$ of the processors ($\rho_{j,i}^t$ is an arbitrary non-negative real number). Let $t_{j,i}$ denote the completion time of job $J_{j,i}$; $t_{j,i}$ is the first date verifying $\int_{a_j + w_{j,i}^s}^{t_{j,i}} \rho_{j,i}^t dt = w_{j,i}^p$. We say that a batch is completed as soon as all its jobs are completed; let t_j denote the completion time of batch B_j , $t_j = \max_{i=1, \dots, u_j} t_{j,i}$. The goal is to minimize the *flow time* of the batches, *i.e.* $\mathcal{B}\text{-FlowTime} = \sum_{B_j \in \mathcal{B}} (t_j - a_j)$. We denote by $\mathcal{B}\text{OPT}_s(\mathcal{B})$ the flow time of an optimal s -speed schedule for \mathcal{B} .

Similarly to the broadcast setting, we say that a request B_j (resp., a job $J_{j,i}$) is *alive* at time t if $a_j \leq t \leq t_j$ (resp., $a_j \leq t \leq t_{j,i}$).

EQUI \circ A Algorithms Family. Given a job scheduling algorithm A , we define the *batches* scheduling algorithm **EQUI \circ A** as follows. Let $R(t)$ denote the set of batches that are alive at time t . **EQUI \circ A** allots to each batch alive at time t an equal amount of processors, *i.e.*, $p/|R(t)|$; then, it runs algorithm A on each alive batch B_j to decide how to split the amount of processors allotted to B_j among its own alive jobs $J_{j,i}$. In the following, we only require algorithm A to be *fully*

active, *i.e.*, that it *allots at all times all the amount of processors* it is given to the alive jobs (*i.e.*, never idles on purpose). Under this requirement, our results hold independently of the choice of A . Examples of fully active algorithms A are: $A = \mathbf{EQUI}$ which equally splits the amount of processors; or $A = \mathbf{MINIDX}$ which allots all the amount of processors to the smallest indexed alive job $J_{j,i}$ in B_j , *i.e.* $i = \min\{i' : J_{j,i'} \text{ is alive at time } t\}$.

Analysis of EQUI \circ A. To analyze the competitiveness of **EQUI \circ A**, we associate to each batches scheduling instance \mathcal{B} , two instances, \mathcal{J}' and \mathcal{J}'' , of job scheduling. We first bound the performances of our algorithm **EQUI \circ A** on \mathcal{B} from above by the performances of **EQUI** on \mathcal{J}' (Lemma 4.1). We then use the “harder” job instance \mathcal{J}'' to show that the job instance \mathcal{J}' was in fact “easier” than the batch instance \mathcal{B} if one increases slightly the number of processors (Lemmas 4.2 and 4.3). Since **EQUI** is competitive on \mathcal{J}' , we can then conclude on the competitiveness of **EQUI \circ A** on \mathcal{B} (Theorem 4.1).

Consider a Seq-Par batches scheduling instance $\mathcal{B} = \{B_1, \dots, B_q\}$ where each batch $B_j = \{J_{j,1}, \dots, J_{j,u_j}\}$ arrives at time a_j and each $J_{j,i}$ in B_j consists of a sequential phase of work $w_{j,i}^s$, followed by a parallel phase of work $w_{j,i}^p$. Consider the s -speed schedule obtained by running algorithm **EQUI \circ A** on instance \mathcal{B} ; let $\rho_{j,i}^t$ denote the amount of processors allotted by **EQUI \circ A** to job $J_{j,i}$ at time t , and $\rho_j^t = \sum_{J_{j,i} \in B_j} \rho_{j,i}^t$ denote the amount of processors allotted to batch B_j at time t ; let $t_{j,i}$ (resp., t_j) be the completion time of job $J_{j,i}$ (resp., batch B_j). We define a Seq-Par job scheduling instance $\mathcal{J}' = \{J'_1, \dots, J'_q\}$, where each job J'_j arrives at time a_j , and is composed of a sequential phase of work $w_j^s = \max_{J_{j,i} \in B_j} w_{j,i}^s$, followed by a parallel phase of work $w_j^p = \int_{a_j + w_j^s}^{t_j} \rho_j^t dt$; intuitively, w_j^s is the length of the longest sequential phase among the jobs in B_j and w_j^p is the total amount of parallel work in B_j to be scheduled by **EQUI \circ A** after the completion of the last sequential phase among the jobs in B_j .

The key to the next lemma is that one gets exactly the same job schedule of the jobs in \mathcal{J}' by running algorithm **EQUI** on instance \mathcal{J}' as by allotting at all times to each job J'_j the same amount of processors as the jobs in B_j received from **EQUI \circ A**.

LEMMA 4.1. (REDUCTION TO JOB SCHEDULING) *If A is fully active, then*

$$\mathbf{EQUI}_s \circ \mathbf{A}(\mathcal{B}) = \mathbf{EQUI}_s(\mathcal{J}').$$

Proof. As long as the longest sequential phase among the jobs in batch B_j is not completed, the batch B_j is alive. By construction, job J'_j is also alive as long as this sequential phase is not completed. Since the amount of processors given to batch B_j in **EQUI \circ A** is given

by **EQUI**, and since **EQUI** is non-clairvoyant, **EQUI**◦**A** allots the same amount of processors to B_j as **EQUI** allots to J'_j until the completion of the longest sequential phase among the jobs in batch B_j . By construction, the longest sequential phase in batch B_j and the sequential phase of J'_j end at the same time and at this moment, all the jobs alive in B_j are in their parallel phase. Thus by construction, the overall amount of remaining parallel work in B_j at that time is equal to the parallel work assigned to J'_j . By construction, the amount of processors given to J'_j equals the amount of processors allotted to batch B_j which is in turn equal to the total amount allotted to each of its remaining alive jobs since A is fully active. The overall remaining amount of parallel work is thus identical in J'_j and B_j until they complete at the same time. Their flow times are thus identical in both schedules. We conclude the proof by reasoning inductively on the completion times (sorted in non-decreasing order) of each phase of each job in each batch. \square

We now define the job instance $\mathcal{J}'' = \{J''_1, \dots, J''_q\}$. \mathcal{J}'' is a kind of worst case instance of the batch instance \mathcal{B} , where all the parallel work in each batch B_j has to be scheduled after the longest sequential phase in B_j . Job J''_j arrives at time a_j and consists of a sequential phase of work $w''_j{}^s = \max_{J_{j,i} \in B_j} w''_{j,i}{}^s$, followed by a parallel phase of work $w''_j{}^p = \sum_{J_{j,i} \in B_j} w''_{j,i}{}^p$.

LEMMA 4.2. (\mathcal{J}' IS EASIER THAN \mathcal{J}'')

$$\text{JOPT}_s(\mathcal{J}') \leq \text{JOPT}_s(\mathcal{J}'').$$

Proof. Since for all j , the sequential works of jobs J'_j and J''_j are identical and the parallel work in J'_j is bounded from above by the parallel work in J''_j , any schedule of \mathcal{J}'' is valid for \mathcal{J}' . \square

LEMMA 4.3. (\mathcal{J}'' WITH δ EXTRA PROCESSORS IS “ALMOST AS EASY” AS \mathcal{B}) For all $\delta > 0$,

$$\text{JOPT}_{1+\delta}(\mathcal{J}'') \leq (1 + 1/\delta) \text{BOPT}_1(\mathcal{B}).$$

Proof. The proof consists in showing that when δ extra processors are given, delaying the completion of each batch B_j by a constant factor, $(1 + 1/\delta)$, allows to postpone the schedule of all the parallel job phases in B_j after the completion of the last sequential phase in B_j , which concludes the proof by construction of \mathcal{J}'' .

Sort the batches of \mathcal{B} by non-increasing arrival time, *i.e.*, assume $a_1 \geq a_2 \geq \dots \geq a_q$. Consider an optimal schedule BOPT_1 of batches B_1, \dots, B_q on one processor. We show by induction that there exists a schedule \mathfrak{S} of \mathcal{J}'' on $1 + \delta$ processors such that each

job J''_j completes before time $t_j + f_j/\delta$, where t_j and $f_j = t_j - a_j$ denote the completion time and the flow time of B_j in BOPT_1 , respectively. We now show that the parallel phase of each job J''_j can be scheduled between time t_j and $t_j + f_j/\delta$; this concludes the proof since, by construction, the sequential phase of J''_j is necessarily completed before t_j . Start with the first job J''_1 . Clearly, $w''_1{}^p \leq f_1$. Thus, the total parallel phase of J''_1 can be scheduled on the δ extra processors between time t_1 and $t_1 + f_1/\delta$. Assume now that the parallel phases of jobs J''_1, \dots, J''_{j-1} have been scheduled in \mathfrak{S} during the time intervals $[t_1, t_1 + f_1/\delta], \dots, [t_{j-1}, t_{j-1} + f_{j-1}/\delta]$ respectively, and consider job J''_j . Since the jobs are considered in non-increasing arrival times, each job J''_k whose parallel phase has been scheduled in \mathfrak{S} between t_j and $t_j + f_j/\delta$ arrived in the time interval $T = [a_j, t_j + f_j/\delta]$ and furthermore $t_k \leq t_j + f_j/\delta$. The total parallel work W of all the jobs currently scheduled in \mathfrak{S} during T , is then in fact scheduled completely in BOPT_1 during T . Note that the parallel work of J''_j was also scheduled in BOPT_1 during this time interval. Since BOPT_1 uses only one processor, we conclude that $W + w''_j{}^p \leq t_j + f_j/\delta - a_j = (1 + 1/\delta)f_j$. As one can schedule up to $(1 + \delta)f_j/\delta = (1 + 1/\delta)f_j$ parallel work between time t_j and $t_j + f_j/\delta$ on $1 + \delta$ processors, the parallel work $w''_j{}^p$ of J''_j can be scheduled in \mathfrak{S} on time. \square

We can now conclude the analysis of **EQUI**◦**A**.

THEOREM 4.1. (COMPETITIVENESS OF **EQUI**◦**A**)

For all $\epsilon > 0$ and $\delta > 0$, **EQUI**◦**A** is a $(2 + \epsilon)(1 + \delta)$ -speed $(2 + 4/\epsilon)(1 + 1/\delta)$ -competitive algorithm for the Non Clairvoyant Seq-Par Batches Scheduling problem.

Proof. We use the result of [7] on the competitiveness of **EQUI** for the non-clairvoyant job scheduling problem to conclude the proof:

EQUI $_{(2+\epsilon)(1+\delta)}$ ◦**A**(\mathcal{B})

$$= \mathbf{EQUI}_{(2+\epsilon)(1+\delta)}(\mathcal{J}') \quad (\text{Lemma 4.1})$$

$$\leq (2 + 4/\epsilon) \text{JOPT}_{(1+\delta)}(\mathcal{J}') \quad (\text{Theorem 1 in [7]})$$

$$\leq (2 + 4/\epsilon) \text{JOPT}_{(1+\delta)}(\mathcal{J}'') \quad (\text{Lemma 4.2})$$

$$\leq (2 + 4/\epsilon)(1 + 1/\delta) \text{BOPT}_1(\mathcal{B}). \quad (\text{Lemma 4.3})$$

\square

5 Competitiveness of **B-EQUISET**

Consider an instance of the online data broadcast problem with dependencies: a set $\mathcal{S} = \{S_1, \dots, S_q\}$ of q requests with arrival times a_1, \dots, a_q , over n items I_1, \dots, I_n of lengths ℓ_1, \dots, ℓ_n . Let \mathcal{E}_s be the s -speed schedule designed by **B-EQUISET** on instance \mathcal{S} , and

$\mathbf{B-EQUISET}_s(\mathcal{S})$ be its flow time. Let \mathcal{O}_1 be a 1-speed optimal schedule of \mathcal{S} , and $\text{BOPT}_1(\mathcal{S})$ be its flow time.

Following the steps of [8], we define an instance \mathcal{B} of non-clairvoyant seq-par batches scheduling from \mathcal{E}_s and \mathcal{O}_1 , such that the performances of $\mathbf{B-EQUISET}$ on \mathcal{S} can be compared to the performances of $\mathbf{EQUIO}A$ on \mathcal{B} for a particular fully-active algorithm A . More precisely, we construct \mathcal{B} such that 1) the flow time of $\mathbf{EQUIO}A$ on \mathcal{B} bounds from above the flow time of $\mathbf{B-EQUISET}$ on \mathcal{S} and 2) the (batches) optimal flow time for \mathcal{B} is at most the (broadcast) optimal flow time for \mathcal{S} if it is given extra resources. Since $\mathbf{EQUIO}A$ is competitive, we can then bound the performances of $\mathbf{B-EQUISET}$ with respect to the (batches) optimal flow time of \mathcal{B} which is by 2) bounded by the (broadcast) optimal flow time of \mathcal{S} .

The intuition behind the construction of \mathcal{B} is the following. A batch of all-new jobs is created for each newly arrived request, with one job per requested item. Each job J stays alive until its corresponding item I is served in \mathcal{E}_s . J is assigned at most two phases depending on the relative service times of I in \mathcal{E}_s and \mathcal{O}_1 . The sequential phase of J lasts until either I is served in \mathcal{E}_s , or the broadcast of I starts in \mathcal{O}_1 . Intuitively, this means that it is useless to assign processors to J before the optimal schedule does. At the end of its sequential phase, if J is still alive, its parallel phase starts and lasts until the broadcast of I is completed in \mathcal{E}_s ; the parallel work for J is thus defined as the total amount of bandwidth that its corresponding item I received within J 's corresponding (broadcast) request in $\mathbf{B-EQUISET}$. By construction, with a suitable choice of A , $\mathbf{EQUIO}A$ constructs the exact same schedule as $\mathbf{B-EQUISET}$ and claim 1) is verified. Concerning claim 2), the key is to consider the jobs corresponding to the broadcast requests for a given item I that are served by a given broadcast of I in \mathcal{O}_1 starting at some time t . The only jobs among them that will receive a parallel phase, are the one for which the broadcast of I in \mathcal{E}_s starts just before or just after t . By construction, the total amount of parallel work assigned to these jobs corresponds to the bandwidth assigned to the two broadcasts of item I by \mathcal{E}_s that start just before and just after time t , each of them being bounded by the length of I . The total amount of parallel work in the jobs for which the broadcast of the corresponding item I starts in \mathcal{O}_1 at some time t , is then bounded by twice the length of I , and can thus be scheduled during the broadcast of I in \mathcal{O}_1 if one doubles the number of processors, which proves claim 2). The following formalizes the reasoning exposed above.

The Job Set Instance \mathcal{B} . Recall the broadcast instance \mathcal{S} , and the two broadcast schedules \mathcal{E}_s and

\mathcal{O}_1 , defined at the beginning of this section, as well as the notations given in Section 2. In particular, let $C_s^\mathcal{E}(I_i, t)$ denote the completion time of the broadcast of item I_i that starts just after t in \mathcal{E}_s , and $B_1^\mathcal{O}(I_i, t)$ be the time of the beginning of the first broadcast of item I_i that starts after t in \mathcal{O}_1 (see Section 2). Recall the description of algorithm $\mathbf{B-EQUISET}$ in Section 3: at time t , let $R(t)$ be the set of alive requests; $\mathbf{B-EQUISET}$ splits equally the bandwidth s among the alive requests and for each alive request S_j , it assigns an arbitrary rate $r_{j,i}(t)$ to each alive item I_i in S_j , such that $\sum_{I_i \text{ alive in } S_j} r_{j,i}(t) = s/|R(t)|$; $\mathbf{B-EQUISET}$ broadcasts then each item I_i at a rate $r_i(t) = \sum_j r_{j,i}(t)$ at time t .

Given \mathcal{S} , \mathcal{E}_s and \mathcal{O}_1 , we define the non-clairvoyant batches scheduling instance $\mathcal{B} = \{B_1, \dots, B_q\}$, where each batch B_j is released at the same time as S_j , *i.e.* at time a_j , and contains one seq-par job $J_{j,i}$ for each item $I_i \in S_j$ (note that the indices i of the jobs $J_{j,i}$ in each batch B_j may not be consecutive depending on the content of S_j). Each job $J_{j,i}$ consists of a sequential phase of work $w_{j,i}^s = (\min\{C_s^\mathcal{E}(I_i, a_j), B_1^\mathcal{O}(I_i, a_j)\} - a_j)$, followed by a parallel phase of work $w_{j,i}^p$. If $C_s^\mathcal{E}(I_i, a_j) \leq B_1^\mathcal{O}(I_i, a_j)$, then $w_{j,i}^p = 0$; otherwise, $w_{j,i}^p = \int_{B_1^\mathcal{O}(I_i, a_j)}^{C_s^\mathcal{E}(I_i, a_j)} r_{j,i}^{\mathcal{E}_s}(t) dt + \eta$ where η is an infinitely small amount of work, *i.e.* if the download of item I_i in request S_j is completed in \mathcal{E}_s after it starts in \mathcal{O}_1 , then the amount of parallel work assigned to $J_{j,i}$ is just slightly higher than the total amount of bandwidth allotted to item I_i within the bandwidth allotted to request S_j by $\mathbf{B-EQUISET}_s$ after the beginning of the corresponding broadcast in \mathcal{O}_1 . Adding an infinitely small amount of work η to the parallel phase of $J_{j,i}$ does not change the optimal batches schedule (except on a negligible (discrete) sets of dates) but since the algorithm $\mathbf{EQUIO}A$ is non-clairvoyant, this ensures that the job $J_{j,i}$ remains alive until the broadcast of item I_i completes even if $\mathbf{B-EQUISET}_s$ deliberately chooses not to broadcast item I_i in the bandwidth allotted to request S_j (the introduction of infinitely small extra load can be rigorously formalized by adding an exponentially decreasing extra load $\gamma/2^k$ to the k th requested job for a small enough γ).

LEMMA 5.1. *There exists a fully-active algorithm A such that: $\mathbf{B-EQUISET}_s(\mathcal{S}) \leq \mathbf{EQUI}_s\mathbf{O}A(\mathcal{B})$.*

Proof. The proof follows the lines of [8]. Given an amount of processors ρ for an alive batch B_j , algorithm A assigns to each alive job $J_{j,i}$ in B_j at time t the same amount of processors as $\mathbf{B-EQUISET}_s$ would have assigned at time t to the corresponding alive item I_i of the corresponding alive request S_j which would have

been assigned a bandwidth ρ . Since **B-EQUISET** _{s} allots all the bandwidth available to alive jobs, A is fully-active. Now, since η is infinitely small, this extra load does not affect the allocation of processors computed by **EQUI** _{s} $\circ A$ except over a negligible (discrete) set of dates. By immediate induction, each job $J_{j,i}$ remains alive in the schedule computed by **EQUI** _{s} $\circ A$, as long as item I_i is alive in batch B_j in \mathcal{E}_s . This is clear as long as $J_{j,i}$ is in its sequential phase. Once $J_{j,i}$ enters its parallel phase, as long as the broadcast of item I_i is not completed, either I_i is broadcast by **B-EQUISET** _{s} in batch B_j and $J_{j,i}$ is scheduled by **EQUI** _{s} $\circ A$ (A copies **B-EQUISET** _{s}), or **B-EQUISET** _{s} deliberately chooses not to broadcast the alive item I_i and since $J_{j,i}$ has an infinitely small amount of extra work, $J_{j,i}$ remains alive in **EQUI** _{s} $\circ A$ as well. The flow time for each job $J_{j,i}$ is then at least the flow time of the corresponding item I_i in \mathcal{E}_s ; we conclude that each batch B_j completes in **EQUI** _{s} $\circ A$ no earlier than its corresponding request S_j in **B-EQUISET** _{s} . \square

LEMMA 5.2. *There exists a 2-speed batches schedule Υ_2 such that: $\Upsilon_2(\mathcal{B}) \leq \text{B-FlowTime}(\mathcal{O}_1)$.*

Proof. Again, the proof follows the lines of [8]. Consider an item I_i . We partition the requests S_j containing item I_i into classes $\mathcal{C}_1, \mathcal{C}_2, \dots$, one for each broadcast of I_i in \mathcal{O}_1 . The k -th class \mathcal{C}_k contains all the requests S_j that download I_i in \mathcal{O}_1 during its k th broadcast, *i.e.* all requests S_j such that $b_1^{\mathcal{O}}(I_i, k-1) < a_j \leq b_1^{\mathcal{O}}(I_i, k)$ (see Section 2 for notations). We show that for all k , the total parallel phases of the jobs $J_{j,i}$ such that $S_j \in \mathcal{C}_k$, can be shoehorned into twice the area of bandwidth allotted by \mathcal{O}_1 to the k th broadcast of item I_i . Since this holds for all i and all k , we obtain a 2-speed schedule Υ_2 such that $\Upsilon_2(\mathcal{B}) \leq \text{B-FlowTime}(\mathcal{O}_1)$.

Let $t_1 = b_1^{\mathcal{O}}(I_i, k)$ be the time of the beginning of the k th broadcast of I_i in \mathcal{O}_1 . Consider a request S_j in class \mathcal{C}_k , clearly $a_j \leq t_1$. By construction, job $J_{j,i}$ is assigned a non-zero parallel work only if S_j completes the download of I_i after t_1 in **B-EQUISET** _{s} . Since S_j arrives before t_1 , it downloads I_i during one of the two broadcasts of I_i in **B-EQUISET** _{s} that start just before or just after t_1 ; let \mathcal{C}_k^- (resp. \mathcal{C}_k^+) be the set of requests served by the broadcast that starts just before t_1 (resp. just after t_1). Let t_2 and t_3 be the completion times of the broadcast of I_i in **B-EQUISET** _{s} that start just before and just after t_1 respectively. By construction, the total amounts W^- and W^+ of parallel work assigned to the jobs $J_{j,i}$ such that $S_j \in \mathcal{C}_k^-$ and

\mathcal{C}_k^+ are respectively:

$$W^- = \sum_{j: S_j \in \mathcal{C}_k^-} \int_{t_1}^{t_2} r_{j,i}(t) dt$$

$$\text{and } W^+ = \sum_{j: S_j \in \mathcal{C}_k^+} \int_{t_1}^{t_3} r_{j,i}(t) dt.$$

Let us rewrite $W^- + W^+ = R_1 + R_2$ with

$$R_1 = \int_{t_1}^{t_2} \sum_{j: S_j \in \mathcal{C}_k} r_{j,i}(t) dt \leq \int_{t_1}^{t_2} r_i(t) dt$$

$$\text{and } R_2 = \int_{t_2}^{t_3} \sum_{j: S_j \in \mathcal{C}_k^+} r_{j,i}(t) dt \leq \int_{t_2}^{t_3} r_i(t) dt.$$

R_1 and R_2 are thus at most the total area allotted to item I_i by **B-EQUISET** _{s} during the broadcasts of I_i that start just before and just after t_1 ; since a broadcast is completed as soon as the rates sum up to the length of the items, $R_1 \leq \ell_i$ and $R_2 \leq \ell_i$, and thus $W^- + W^+ \leq 2\ell_i$. Since \mathcal{O}_1 allots a total bandwidth of ℓ_i to broadcast item I_i after time t_1 , and since the parallel works of the jobs $J_{j,i}$ such that $S_j \in \mathcal{C}_k$ are released at time t_1 and sum up to a total $W^- + W^+ \leq 2\ell_i$, one can construct on 2 processors, a 2-speed schedule Υ_2 in which the parallel phases of each of these jobs $J_{j,i}$ completes before the k th broadcast of I_i completes in \mathcal{O}_1 .

Since no processor needs to be allotted to the sequential phases, repeating the construction for each item I_i yields a valid 2-speed schedule Υ_2 in which each job $J_{j,i}$ completes before the corresponding request S_j completes the download of I_i in \mathcal{O}_1 . It follows that each batch B_j is completed in Υ_2 before its corresponding request S_j is served by \mathcal{O}_1 . \square

We now conclude with the proof of the main theorem.

Proof of Theorem 1.1. Setting $s = (4 + \epsilon)(1 + \delta)$, the competitiveness of **EQUI** $\circ A$ (Theorem 4.1) concludes the result:

$$\begin{aligned} & \mathbf{B-EQUISET}_{(4+\epsilon)(1+\delta)}(\mathcal{S}) \\ & \leq \mathbf{EQUI}_{(4+\epsilon)(1+\delta)} \circ \mathbf{A}(\mathcal{B}) \quad (\text{Lemma 5.1}) \\ & \leq (2 + 8/\epsilon)(1 + 1/\delta) \mathbf{BOPT}_2(\mathcal{B}) \quad (\text{Theorem 4.1}) \\ & \leq (2 + 8/\epsilon)(1 + 1/\delta) \Upsilon_2(\mathcal{B}) \\ & \leq (2 + 8/\epsilon)(1 + 1/\delta) \mathbf{BOPT}_1(\mathcal{S}) \quad (\text{Lemma 5.2}) \end{aligned}$$

\square

The B-EQUISET-EDF algorithm. We apply the same method as in [8]. Let $s = (4 + \epsilon)(1 + \delta)^2$ and $c = (2 + 8/\epsilon)(1 + 1/\delta)^2$. **B-EQUISET-EDF** simulates the $s/(1 + \delta)$ -speed execution of **B-EQUISET** and at each time t such that the broadcast of an item I_i in

B-EQUISET is completed, it releases an item I'_i of length ℓ_i with a deadline $t + (t - t')/\delta$ where t' is the time of the beginning of the considered broadcast of I_i in **B-EQUISET**. Then, **B-EQUISET-EDF** schedules on a bandwidth s each item I'_i according the earliest-deadline-first policy. With an argument similar to Lemma 4.3 or [8], one can show that a feasible schedule of the items I'_i exists and thus that earliest-deadline-first constructs it which ensures that **B-EQUISET-EDF** is s -speed c -competitive. Since earliest-deadline-first preempts the broadcast of an item only when a new item arrives, **B-EQUISET-EDF** preempts each broadcast at most once on average. Note that one can avoid long idle period in **B-EQUISET-EDF**'s schedule by broadcasting an arbitrary item I_i alive in **B-EQUISET** at time t if no item I'_i is currently alive.

6 Concluding remarks

Several directions are possible to extend this work. First, **B-EQUISET** does not have precise policy to decide in which order one should broadcast the items within each requested set; deciding on a particular policy may lead to better performances (bandwidth and/or competitive ratio). Second, it might be interesting to design a longest-wait-first greedy algorithm in presence of dependencies; **B-EQUISET** shows that the items should not simply receive bandwidth according to the number of outstanding requested sets for this item (the allotted bandwidth depends also on the number of outstanding items within each outstanding set), it is thus a challenging question to design proper weights to aggregate the current waits of the requested sets including a given item.

References

- [1] N. Bansal, M. Charikar, S. Khanna, and J. S. Naor. Approximating the average response time in broadcast scheduling. In *Proc. of the 16th ACM-SIAM Symp. on Discrete Algorithms (SODA 2005)*, pages 344–353, 2005.
- [2] N. Bansal, D. Coppersmith, and M. Sviridenko. Improved approximation algorithms for broadcast scheduling. In *Proc. of the An. ACM-SIAM Symp. on Discrete Algorithms (SODA 2006)*, pages 344–353, 2006.
- [3] A. Bar-Noy, J. S. Naor, and B. Schieber. Pushing dependent data in clients-providers-servers systems. *Wireless Networks*, 9:421–230, 2003.
- [4] A. Bar-Noy and Y. Shilo. Optimal broadcasting of two files over an asymmetric channel. *J. Parallel Distrib. Comput.*, 60(4):474–493, 2000. (also published in the Proc. of INFOCOM 1999).
- [5] K. Cai, H. Lin, and C. Chen. Correlation-based data broadcasting in wireless networks. In *LNCS Proc. of the 2nd British National Conf. on Databases (BNCOD 22)*, volume 3567, pages 111–119, Jul. 2005.
- [6] S. Dey and N. Schabanel. Customized newspaper broadcast: Data broadcast with dependencies. In *Proc. of Lat. Am. Conf. on Theor. INformatic (LATIN)*, pages 362–373, 2006.
- [7] J. Edmonds. Scheduling in the dark. In *STOC*, pages 179–188, 1999.
- [8] J. Edmonds and K. Pruhs. Multicast pull scheduling: When fairness is fine. *Algorithmica*, 36(3):315–330, 2003. (also published in Proc. of SODA 2002).
- [9] J. Edmonds and K. Pruhs. A maiden analysis of longest wait first. *ACM Trans. Algorithms*, 1(1):14–32, 2005.
- [10] T. Erlebach and A. Hall. NP-hardness of broadcast scheduling and inapproximability of single-source unsplittable min-cost flow. In *Proc. of Symp. on Discrete Algorithms (SODA)*, pages 194–202, 2002.
- [11] R. Gandhi, S. Khuller, Y. A. Kim, and Y.-C. J. Wan. Algorithms for minimizing response time in broadcast scheduling. In *Proc. of Conf. on Int. Prog. and Comb. Opt. (IPCO)*, pages 425–438, 2002.
- [12] J.-L. Huang and M.-S. Chen. Broadcast program generation for unordered queries with data replication. In *Proc. of the ACM Symp. on Applied Computing (SAC 2003)*, pages 866–870, Mar. 2003.
- [13] J.-L. Huang and M.-S. Chen. Dependent data broadcasting for unordered queries in a multiple channel mobile environment. *IEEE Trans. on Knowledge and Data Engineering*, 16(9):1143–1156, Sep. 2004.
- [14] B. Kalyanasundaram, K. Pruhs, and M. Velauthapillai. Scheduling broadcasts in wireless networks. In *Proc. of European Symp. on Algorithms*, pages 290–301, 2000.
- [15] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge university press, 1995.
- [16] A. Yao. Probabilistic computations: Towards a unified measure of complexity. In *Proc. of 17th Symp. on Fond. of Computer. Science (FOCS)*, pages 222–227, 1977.