The Data Broadcast Problem with Non-Uniform Transmission Times (Revised version)

Claire Kenyon* Nicolas Schabanel†

October 9, 2001 at 11:18

Abstract. The data broadcast problem consists in finding an infinite schedule to broadcast a given set of messages so as to minimize the average service time to clients requesting messages, and the cost of the broadcast. This problem also models the maintenance scheduling problem and the multi-item replenishment problem.
Previous work concentrated on a discrete-time restriction where all messages have transmission time equal to 1. Here, we study a generalization of the model to the more realistic setting of continuous time and messages of non-uniform transmission times. We prove that the data broadcast problem is *NP*-hard, even if the broadcast costs are all zero, and give 3-approximation algorithms for broadcasting messages on a single channel.
This is the full version of [KS99].

Keywords: Approximation Algorithms, Scheduling, Randomized Algorithms, NP-completeness, Data Broadcasting.

Contents

*LRI, Université Paris XI, Bâtiment 490, F-91405 Orsay, France, http://www.lri.fr/~kenyon
†LIP, ENS Lyon, 46, allée d'Italie, F-69364 Lyon Cedex 07, France, http://www.ens-lyon.fr/~nschaban

1 Introduction

1.1 Motivation

This paper studies an optimization problem which arises in three contexts: Data broadcasting, Scheduling maintenance service, and Multi-item replenishment.

Broadcasting is an efficient means of disseminating data in wireless communication environments, where there is a much larger communication capacity from the information source to the recipients than in the reverse direction, such as when mobile clients retrieve information from a server base-station through a wireless medium. In broadcasting protocols, items are scheduled continously, over an infinite time-horizon. The requests do not propagate in the system, but the clients wait for the requested items to be broadcast, thus making the system *pseudointeractive* or *push-based* (as opposed to *pull-based*): the schedule is independent of the incoming requests, since it is oblivious to them. Push-based systems allow to save bandwidth and to reduce service time by unloading the server from treating the request and by reducing the bandwidth occupied by the most popular items [Bes96, AFZ97]. [Ach98, Sch00a] present a very complete history of the field. An efficient broadcast scheme seeks to minimize the average service time, (*i.e.* the amount of time spent by a mobile client to obtain a desired piece of information in the broadcast), while also minimizing the resulting broadcast cost (*e.g.* in the context of data transfers on the World Wide Web, each message has a broadcast cost, the required bandwidth, which is proportional to its length). While most previous work made the simplifying assumption that all items had the same transmission time, our main focus in this paper is to deal rigorously with non-uniform transmission times.

The maintenance service problem schedules m machines for maintenance over an infinite time horizon and seeks to minimize both the costs associated with each maintenance and the operation costs of the machines, where the operation costs of the machines are assumed to increase with the time elapsed since the last maintenance. Here again it seems reasonable to consider the case where sophisticated machines have longer maintenance time than others.

The multi-item replenishment problem consists in, given m items types, deciding over time which item stock needs to be replenished given holding costs, ordering costs, and the rate at which each item is consumed.

All three problems can be modeled similarly.

1.2 Problem definition

In this paper we will adopt the Data-broadcast terminology. The input consists of m messages M_i, $i = 1 \dots m$, defined by their *lengths* ℓ_i (the time required to broadcast M_i), their *request probabilities* (or *popularities*) p_i, and their *broadcast cost* c_i. The problem is to decide in what sequence S to schedule the broadcast messages over an infinite time-horizon, so as to minimize the sum of the average service time to Poisson requests and of the average broadcast cost, i.e. so as to minimize $\limsup_{T\to\infty} \{\mathcal{A}\mathrm{ST}(S,[0,T]) + \mathcal{A}\mathrm{BC}(S,[0,T])\}$;

here $\mathcal{A}\mathrm{ST}(S,[0,T])$ denotes the average service time to a request which *1)* is generated at a random uniform instant between 0 and T, *2)* requests message M_i with probability p_i, and *3)* must then wait until the start of the next broadcast of M_i to *4)* start downloading M_i (which takes ℓ_i time); and $\mathcal{A}\mathrm{BC}(S,[0,T])$ is the average broadcast cost of the messages whose broadcast starts between the dates 0 and T. As remarked above, the broadcast scheme is pseudointeractive: the actual requests do not propagate beyond the client, and the schedule is oblivious to the actual sequence of requests. Note that when this definition is specialized to the uniform-lengths setting, our definition agrees with the literature on data broadcasting (see Lemma 1).

In the Maintenance Scheduling problem [AGH95, BNBNS98], the input consists of m machines M_i, $i = 1..m$, defined by the *length* ℓ_i *of their maintenance*, their *operation cost rate* p_i, and their *maintenance cost* c_i. The instantaneous operation cost of a machine is an affine function of the time elapsed since its last maintenance, with linear rate p_i and constant $b \cdot p_i$ (the cumulated cost since the last maintenance thus increases quadratically). The problem is to find an maintenance schedule over an infinite time-horizon, so as to minimize the total costs incurred, i.e. the sum of the operation costs and of the maintenance cost. This problem is, as we will see in Proposition 3, identical to the Data Broadcast problem. The multi-item replenishment problem consists of, given m items types, deciding over time when to fill which item stock needs given holding costs, ordering costs, and the rate at which each item is consumed. The Multi-item Replenishment problem is a special case of the maintenance scheduling problem (see [BNBNS98, AGH95]).

1.3 Background

As far as we know, almost all previous work focused on the uniform-lengths and discrete-time model, when all messages (or machines) have length (or operating time) equal to 1.

Let us first review results on discrete-time and uniform-lengths Data Broadcast problem without broadcast costs. This problem was first studied in the context of Teletext. In [Gec83], Gecsei analyzes the mean response time of memoryless randomized algorithms and proves that any optimal algorithm from this class should choose the next message to broadcast with probability proportional to the square root of its popularity. In [AW85], Ammar and Wong study periodic broadcast schedules. They derive an algebraic expression for the average service time (which is essentially Lemma 1 below), and prove a lower bound (essentially the one in Lemma 6 below), from which they get inspiration to derive a greedy-type algorithm. They provide numerical evidence that their algorithm is efficient by studying it when the message popularities follow the Zipf distribution, which is claimed to closely approximate real user behavior. In [AW87], they define the mean response time of arbitrary (not necessarily periodic) schedules, analyze structural properties of optimal schedules and prove that there exists an optimal schedule which is periodic. From this, they deduce a finite-time algorithm for computing the optimal solution. They also present an algorithm for constructing a good schedule, which uses the golden ratio sequence [AW87] (They do not provide a performance ratio analysis). Further work [Amm87] attacks a model where successive requests are not independent and prefetching may be used. Greedy heuristics for the same problem are empirically tested in [ST97, AGH95]. This study on the single-channel uniform-length case, was pursued by [AGH95], who obtained a deterministic 2.5-approximation algorithm, and moreover by [BNBNS98] who obtained a deterministic 2-approximation algorithm and proved that the golden ratio heuristic of [AW87] has approximation ratio 9/8. In [KSY00], the authors improve this result by giving a polynomial time approximation scheme ($(1+\varepsilon)$-approximation for any $\varepsilon > 0$) for this uniform length case. Finally in [KZ98], the authors study indexed data broadcast where the objective function to minimize is a combination of response time and tuning time. Note that in the middle 1990's a simplified model was introduced by [AAFZ95]: the broadcast disks. In this model, the goal is to partition the set of messages. Each subset is seen as a rotating disk. On each disk the messages are scheduled in Round Robin order. The rotation speed of each disk is determined according to the popularities of its items. The schedules are obtained by merging the output of the disks. This restrictive setting allows to simplify the model to study prefechting and caching strategies [Ach98] but prevents from approaching the optimal strategy.

Even in the uniform length model, very few articles consider a multi-channel variant of the problem where upto W (the number of channels) messages can be broadcast at the same time. Previous work for the general uniform-length Maintenance Scheduling problem can be found in [BNBNS98]. There, the authors prove that there is an optimal schedule which is periodic. Using the operation costs c_i, they prove that the Maintenance Scheduling Problem is *NP*-hard (but are unable to prove the *NP*-hardness of the Data Broadcast problem). They design a 9/8-approximation algorithms for the Data Broadcast and the Maintenance Scheduling problems, even when there is more than one channel.

[VH97a, VH97b] are the only references which considers non-uniform length messages. They propose a first lower bound (basically our LB_0 in Section 4.1) and report experimental results for several heuristics on one or two channels in comparaison to this lower bound.

1.4 The results

This paper focuses on the case where the message lengths (transmission times) are not equal. This generalization is significantly harder: in [AW87], the authors justify their uniform-length assumption in those terms: "This [uniform lengths] assumption [...] is required to render the problem under consideration tractable". In fact, there are some important structural differences between our model and the uniform-length model: for example, in our model, an optimal schedule need no longer be periodic (see Section 7). In general it is not even clear *a priori* whether an optimal solution can be described in finite time! Thus we try to provide a careful definition and treatment of this model. After proving that periodic schedules are arbitrarily close to optimal, our results are then obtained "by density" of the periodic schedules.

We show that even when there are no broadcast costs, the problem is *NP*-hard (Note that in [BNBNS98] *NP*-hardness is proved only when broadcast costs are present).

Then, we present a polynomial 3-approximation for the problem. To design that algorithm, we first observe that, as stated by the authors of [BNBNS98], the natural extension of the lower bound of [AW85] to our model is no longer tight in the non-uniform-length case, and can, in fact, be arbitrarily bad and can yield to arbitraly bad approximations. One important ingredient of our algorithm here consists in designing a second lower bound which is tight up to a constant factor. This is the key to finding a constant-factor approximation. Our approximation outputs a periodic schedule with quadratic period. It relies on a randomized algorithm that is derandomized.

At the end of the paper we introduce a multi-channel framework. Very little is known on this extension. This last section presents only a tour by examples of some of the difference with the single channel case and with the multiple-channel case with uniform transmission time.

1.5 Organization of the paper

The paper is organized as follows. Section 2 presents the notations and facts about the cost of a schedule. It also shows that our definition of the cost function is equivalent to the definition from [AW85, AW87, Amm87] and [AGH95, BNBNS98]. In Section 3, we prove that the cost of periodic schedules is arbitraly close to the optimal cost, allowing us to manipulate only periodic schedules. Section 4 presents our lower bounds. Section 5 proves that the problem is strongly *NP*-hard even if zero broadcast cost are assumed. We give in Section 6 a randomized and a deterministic 3-approximation algorithm. Section 7 finally presents a multiple channel franework where several message can be broadcast. The last section concludes by giving some axes of further research on that subject.

2 Definitions and Notations

2.1 The problem

The input. The messages M_i, $i = 1, \dots, m$, are defined by their *lengths* $\ell_i > 0$, their *request probabilities* or *popularities* $p_i > 0$ such that $p_1 + \cdots + p_m = 1$, and their *broadcast costs* $c_i \geqslant 0$.

The schedule. A *schedule* S of $M_1, \dots, M_m$ is an infinite sequence $M_{s(1)} M_{s(2)} \dots$, where $1 \leqslant s(i) \leqslant m$. We denote by $t(n)$ the starting time of the broadcast of $M_{s(n)}$, which must satisfy $t(n+1) \geqslant t(n) + \ell_{t(n)}$.

A schedule S is *periodic* with *period* $T > 0$, if there exists $N > 0$ such that for any $n \geqslant 1$, $s(n+N) = s(n)$ and $t(n+N) = t(n) + T$. Periodic schedules are interesting because we only need to store the first T unit of times to store the complete schedule.

The objective function. We are interested in minimizing a combination of two quantities on S. The first one, denoted by $\mathcal{A}\mathrm{ST}(S)$, is the *average service time* to a random request (where the average is taken over the moments when requests occur *and* over the type M_i of message requested). If we define for a time interval I, $\mathcal{A}\mathrm{ST}(S, I)$ as the average service time of a random request arriving in I, $\mathcal{A}\mathrm{ST}(S)$ is defined as:

$$\mathcal{A}\mathrm{ST}(S) = \limsup_{T \to \infty} \mathcal{A}\mathrm{ST}(S, [0, T])$$

The second quantity is the *average broadcast cost* $\mathcal{A}\mathrm{BC}(S)$ of the messages, defined as the asymptotic value of the average broadcast cost of S over any time interval (where the average is over time). By definition, each broadcast of a message M_i costs c_i. For a time interval I, the average broadcast cost of S over I, $\mathcal{A}\mathrm{BC}(S, I)$, is defined as the sum of the cost of all the messages whose broadcast begins in I, divided by the length of I. $\mathcal{A}\mathrm{BC}(S)$ is then defined as:

$$\mathcal{A}\mathrm{BC}(S) = \limsup_{T \to \infty} \mathcal{A}\mathrm{BC}(S, [T_0, T]), \quad \text{for any } T_0$$

The quantity which we want to minimize is the *cost* of a schedule S which we define as follows:

$$\mathrm{COST}(S) = \mathcal{A}\mathrm{ST}(S) + \mathcal{A}\mathrm{BC}(S)$$

General notations. We adopt the following notations for the rest of the paper:

- $x_{\max} =_{\mathrm{def}} \max_i x_i$ is the maximum value of the x_is (idem $x_{\min} =_{\mathrm{def}} \min_i x_i$).
- $\mathbb{E}_x[y] =_{\mathrm{def}} \sum_i x_i y_i$ is the expected value the y_is given the discrete distribution (x_i). *E.g.* $\mathbb{E}_p[\ell]$ is the average downloading time since message M_i is requested (and downloaded) with probability p_i.
- $\mathcal{L} =_{\mathrm{def}} \ell_1 + \cdots + \ell_m$ is the sum of the lengths of the messages and $\mathcal{C} =_{\mathrm{def}} c_1 + \cdots + c_m$ is the sum of the broadcast costs of the messages.
- OPT is the optimal cost.

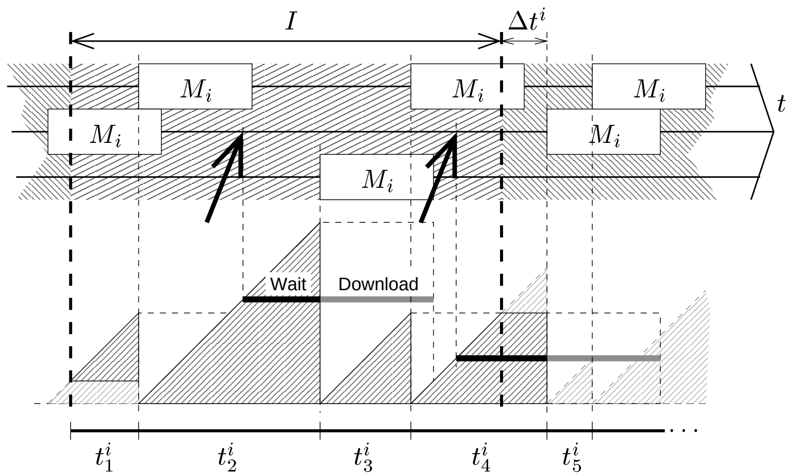

Figure 1: Illustration of the notations and expression in Lemma 1. Here, $n_i = 3$.

2.2 An algebraic expression for the cost

Lemma 1 introduces notations which will be used throughout the paper. Proposition 3 relates our definition of cost to the algebraic definition given in previous work [AW85, BNBNS98] for discrete time and uniform lengths on either data-broadcasting or maintenance scheduling.

Lemma 1 (Algebraic definition of COST**)** *Consider a schedule S of m messages $M_1, \ldots, M_m$ and a time interval I. Take a request generated at a random uniform instant of I and asking for M_i with probability p_i. Then the average service time to the request is:*

$$\mathcal{A}\mathrm{ST}(S, I) = \mathbb{E}_p[\ell] + \frac{1}{2|I|}\sum_{i=1}^{m} p_i \left\{ \sum_{j=1}^{n_i+1} (t_j^i)^2 - (\Delta t^i)^2 \right\}$$

and the average broadcast cost of S over I is:

$$\mathcal{A}\mathrm{BC}(S, I) = \frac{1}{|I|}\sum_{i=1}^{m} n_i c_i$$

where (cf. Figure 1):

- *n_i is the number of broadcasts of M_i starting in I according to S.*
- *t_1^i is the time elapsed from the beginning of I to the beginning of the first broadcast of M_i; and t_j^i, for $j \geqslant 2$, denotes the time elapsed from the $(j-1)$th to the jth broadcast of M_i since the beginning of I.*
- *Δt^i is the length of the interval from the end of I to the first broadcast of M_i starting after I.*

Proof. Consider a request Q in I. Q requests for M_i with probability p_i. With probability $t_j^i/|I|$, for $1 \leqslant j \leqslant n_i$, Q arrives in an interval of length t_j^i and waits $t_j^i/2$ on average until the next broadcast of M_i and then ℓ_i more until the end of the downloading. With probability $(t_{n_i+1}^i - \Delta t^i)/|I|$, Q arrives in the last interval and waits $\left(\frac{t_{n_i+1}^i - \Delta t^i}{2} + \Delta t^i\right) = \frac{t_{n_i+1}^i + \Delta t^i}{2}$ on average until the next broadcast of M_i, plus ℓ_i more until the end of the downloading. Summing over all the intervals and messages yields the expression given for $\mathcal{A}\mathrm{ST}(S, I)$.

The average cost for broadcasting M_i during I is $n_i c_i/|I|$. Summing over all the messages yields the expression given for $\mathcal{A}\mathrm{BC}(S)$. $\square$

Corollary 2 *If S is a periodic schedule of period T, then:*

$$\mathcal{A}\mathrm{ST}(S) = \mathbb{E}_p[\ell] + \frac{1}{2T}\sum_{i=1}^{m} p_i \sum_{j=1}^{n_i} (t_j^i)^2 \text{ and } \mathcal{A}\mathrm{BC}(S) = \frac{1}{T}\sum_{i=1}^{m} n_i c_i$$

where t_j^i denotes the time elapsed from the $(j-1)$st to the jth broadcast of M_i in a period of S.

Proof. For $\mathcal{A}\mathrm{BC}(S)$, the statement is clear. To calculate the average service time, consider requests to M_i. We take a partition of time into intervals of length T, such that *each interval starts exactly at the beginning of a broadcast of M_i*. With probability $t_j^i/|I|$, for $1 \leqslant j \leqslant m$, Q arrives in an interval of length t_j^i and waits $t_j^i/2$ on average until the next broadcast of M_i and then ℓ_i more until the end of the downloading. □

2.3 An equivalent definition of the cost

This section establishes that the problem studied here is the same as in [BNBNS98, AGH95] by showing that the service time is equivalently defined by the average time elapsed since the last broadcast, plus the message length. In other terms, we show here that time is reversible for this problem. This is straightforward for periodic schedules but requires some work for general, non periodic schedules.

Definition. For any given schedule S and any time t, we define the *penalty* $\mathrm{PE}(S,t,M_i)$ of a schedule S at time t for message M_i as p_i multiplied by the time elapsed since the beginning of the last broadcast of M_i. Note that the penalty is proportional to the expected number of unserved requests for M_i at time t.
If t_i is the date of the beginning of the last broadcast of M_i before time t, we have: $\mathrm{PE}(S,t,M_i) = p_i(t-t_i)$ (we assume that no request arrives before time $t=0$). We define then the *average penalty* $\mathcal{A}\mathrm{PE}(S,I)$ *of S over a time interval I* as follows:

$$\mathcal{A}\mathrm{PE}(S,I) = \frac{1}{|I|}\int_I \sum_{i=1}^{m} \mathrm{PE}(S,t,M_i)\, dt$$

and the *average penalty* of S as:

$$\mathcal{A}\mathrm{PE}(S) = \limsup_{T\to\infty} \mathcal{A}\mathrm{PE}(S,[T_0,T]), \quad \text{for any given } T_0$$

We define then a new cost $\widehat{\mathrm{COST}}$ for schedule S as:

$$\widehat{\mathrm{COST}}(S) = \mathbb{E}_p[\ell] + \mathcal{A}\mathrm{PE}(S) + \mathcal{A}\mathrm{BC}(S) \tag{1}$$

This definition of the cost matches the one in [BNBNS98, AGH95]. The next claim proves that the two cost functions COST and $\widehat{\mathrm{COST}}$ are identical. (Note that this property is clear for periodic schedules, since in that case both quantities (see Corollary 2) sum the squares of the interbroadcast lengths.)

Proposition 3 (Time reversibility) *For any schedule S:*

$$\mathcal{A}\mathrm{ST}(S) = \mathcal{A}\mathrm{PE}(S) + \mathbb{E}_p[\ell]$$

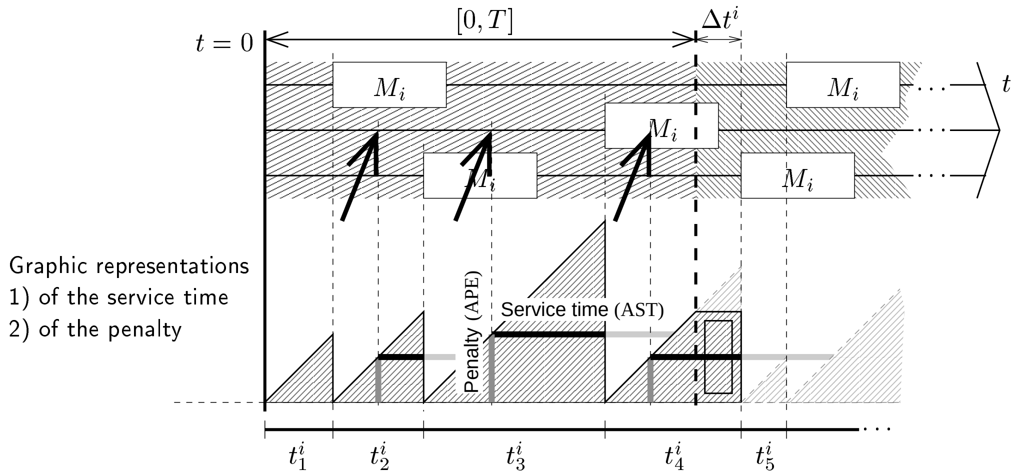

Figure 2: Illustration of Proposition 3. Here, $n_i(T) = 3$.

Proof. Consider a schedule S. We adopt the following notations:

- $n_i(T)$ is the number of broadcasts of M_i that begin in $[0, T]$ in S.
- t_i^j is the time elapsed between the $(j-1)^{\text{th}}$ and j^{th} broadcasts of M_i (we assume that a 0^{th} broadcast of M_i occurs at time $t = 0$).

According to Lemma 1 and also Figure 2, we have:

$$\mathbb{E}_p[\ell] + \frac{1}{2T}\sum_{i=1}^{m} p_i \sum_{j=1}^{n_i(T)} (t_j^i)^2 \leqslant \qquad \mathcal{A}\mathrm{ST}(S,[0,T]) \leqslant \mathbb{E}_p[\ell] + \frac{1}{2T}\sum_{i=1}^{m} p_i \sum_{j=1}^{n_i(T)+1} (t_j^i)^2$$

and,

$$\underbrace{\mathbb{E}_p[\ell] + \frac{1}{2T}\sum_{i=1}^{m} p_i \sum_{j=1}^{n_i(T)} (t_j^i)^2}_{=_{\text{def}}\ A(T)} \leqslant \quad \mathbb{E}_p[\ell] + \mathcal{A}\mathrm{PE}(S,[0,T]) \leqslant \underbrace{\mathbb{E}_p[\ell] + \frac{1}{2T}\sum_{i=1}^{m} p_i \sum_{j=1}^{n_i(T)+1} (t_j^i)^2}_{=_{\text{def}}\ B(T)}$$

We will prove that the lower and upper bounds, $A(T)$ and $B(T)$, above both have the same limsup. This will immediately imply that $\mathcal{A}\mathrm{ST}(S) = \mathcal{A}\mathrm{PE}(S) + \mathbb{E}_p[\ell]$.

If the limsup of $A(T)$ is $+\infty$, then clearly $\mathcal{A}\mathrm{ST}(S) = \mathcal{A}\mathrm{PE}(S) = +\infty$.

Otherwise, take any κ such that for all T large enough, $A(T) \leqslant \kappa$. In particular, for every i, take $A(T_i)$ with $T_i = t_1^i + \cdots + t_{n_i(T)+1} \leqslant T + t_{n_i(T)+1}$, we have

$$\frac{p_i}{2(t_1^i + \cdots + t_{n_i(T)+1}^i)} \sum_{j=1}^{n_i(T)+1} (t_j^i)^2 \leqslant \kappa$$

This implies two things. First, that $(t_{n_i(T)+1}^i)^2 \leqslant \frac{2\kappa}{p_i}(T + t_{n_i(T)+1}^i)$, hence $t_{n_i(T)+1}^i = O(\sqrt{T})$. Second that

$$B(T) = \mathbb{E}_p[\ell] + \frac{1}{2T}\sum_{i=1}^{m} p_i \sum_{j=1}^{n_i(T)+1} (t_j^i)^2 \leqslant \kappa\, \frac{T + \max\limits_{1\leqslant i\leqslant m} t_{n_i(T)+1}^i}{T} \leqslant \kappa + O\left(\frac{1}{\sqrt{T}}\right)$$

This is true for any κ such that $A(T) \leqslant \kappa$ for all T large enough, it follows that the lower and upper bounds, $A(T)$ and $B(T)$, have the same limsup, and so $\mathcal{A}\mathrm{ST}(S) = \mathcal{A}\mathrm{PE}(S) + \mathbb{E}_p[\ell]$. $\square$

3 A reduction to periodic schedules

Previous work relied on the existence of a periodic optimal schedule. Although this statement still holds when there are no broadcast costs (see Lemma 11), we do not currently know whether there is an optimal schedule which is periodic in the case of broadcast costs. However, for the purpose of approximation, it is sufficient to consider periodic schedules, as the following lemma shows.

Lemma 4 (Density of periodic schedules) *Let S be a schedule of $M_1, \ldots, M_m$ on W channels. For any $\varepsilon > 0$, there exists a* periodic *schedule S' whose cost satisfies:*

$$\mathrm{COST}(S') \leqslant \mathrm{COST}(S) + \varepsilon$$

Proof. Take T so that for any $t \geqslant T$: $\mathrm{COST}(S, [0, t]) \leqslant \mathrm{COST}(S) + \varepsilon/2$. Let us choose $t \geqslant T$, matching the date of the end of the broadcast of some message in S, that will be fixed later. Let S_t be the periodic schedule with period $(t + \mathcal{L})$ which broadcasts the same messages as S during $[0, t]$, and then broadcasts during $[t, t + \mathcal{L}]$ the m messages in order of next appearance in S after time t. Thus the requests which arise in S_t during $[0, t]$ are served no later than in S.

A request that arises during $[t, t + \mathcal{L}]$ is served after at most $(\mathcal{L} + \mathrm{ST}(S, 0))$ time, where $\mathrm{ST}(S, 0)$ denotes the average service time for a request that arises at time 0 in S. Thus,

$$\mathcal{A}\mathrm{ST}(S_t) \leqslant \mathcal{A}\mathrm{ST}(S, [0, t]) + \frac{\mathcal{L}(\mathcal{L} + \mathrm{ST}(S, 0))}{t + \mathcal{L}} = \mathcal{A}\mathrm{ST}(S, [0, t]) + O\Big(\frac{1}{t}\Big).$$

As for the broadcast costs, we have:

$$\mathcal{A}\mathrm{BC}(S_t) \leqslant \mathcal{A}\mathrm{BC}(S, [0, t]) + \frac{\mathcal{C}}{t + \mathcal{L}} = \mathcal{A}\mathrm{BC}(S, [0, t]) + O\Big(\frac{1}{t}\Big).$$

Hence:

$$\mathrm{COST}(S_t) \leqslant \mathrm{COST}(S, [0, t]) + O\Big(\frac{1}{t}\Big) \leqslant \mathrm{COST}(S) + \frac{\varepsilon}{2} + O\Big(\frac{1}{t}\Big),$$

and choosing t large enough (so that $O(\frac{1}{t}) \leqslant \frac{\varepsilon}{2}$) yields the result. □

4 Lower bounds

4.1 Extension of known lower bounds

Finding good lower bounds is a key point to designing provably efficient algorithms. The first lower bound LB_0 below is a natural generalization of the one from [VH97a, BNBNS98]. Unfortunately, this lower bound is not tight when the messages have different lengths (see Fact 7). This lower bound will however be used in the *NP*-hardness proof of Theorem 12 and some other examples, which is why we present it in detail. We then design a second, tight lower bound.

Let us define the quantity LB_0 as the following minimization problem:

$$\mathrm{LB}_0 \left\{ \begin{array}{c} \mathbb{E}_p[\ell] \; + \; \min\limits_{\tau > 0} \; \sum\limits_{i=1}^{m} \Big(p_i \frac{\tau_i}{2} + \frac{c_i}{\tau_i} \Big) \\ \text{subject to: } \sum\limits_{i} \frac{\ell_i}{\tau_i} \leqslant 1 \end{array} \right.$$

Lemma 6 will show that it is a lower bound for our problem. But first let us see how to solve this minimization problem.

Note that when all messages have zero broadcast cost, Lemma 5 gives the "square-root rule" in [VH97a] inspired from the uniform length case in [AW85].

Lemma 5 *The minimization problem* LB_0 *admits a unique solution* τ^* *verifying:*

$$\tau_i^* = \sqrt{\frac{2c_i + \lambda^* \ell_i}{p_i}}$$

for a certain $\lambda^* \geqslant 0$ *(independent of* i*), such that:* $\lambda^* = 0$*, if* $\sum_{i=1}^m \ell_i \sqrt{\frac{p_i}{2c_i}} \leqslant 1$*; otherwise,* λ^* *is the unique positive solution to* $\sum_{i=1}^m \ell_i \sqrt{\frac{p_i}{2c_i + \lambda^* \ell_i}} = 1$.

Square-Root Rule: *If all the messages have zero broadcast cost* ($c_i = 0$*, for all* i*), then:*

$$\mathrm{LB}_0 = \mathbb{E}_p[\ell] + \frac{1}{2}\left(\sum_{i=1}^m \sqrt{p_i \ell_i}\right)^2 \quad \text{and} \quad \tau_i^* = \sqrt{\frac{\ell_i}{p_i}\lambda^*}$$

with $\lambda^* = \left(\sum_{j=1}^m \sqrt{p_j \ell_j}\right)^2$.

Proof. As in [BNBNS98] we use Lagrangian relaxation. We explain this part in some detail to justify rigorously the validity of the Lagrangian relaxation technique in this setting. Let $f(\tau) = \sum_{i=1}^m p_i \tau_i + \frac{2c_i}{\tau_i}$ denote the variable term of the objective function, and $g(\tau) = 1 - \sum_{i=1}^m \frac{\ell_i}{\tau_i}$ denote the (main) constraint of the system. The constraints define a subdomain D of $\mathbb{R}_+^m$. There are two cases:

- Either the minimum of f in D is reached in the interior of D, in which case it is a critical point[1] of f. However we note that the particular f in our setting has a unique critical point $\tau^\#$ in $\mathbb{R}_+^m$, defined by $\tau_i^\# = \sqrt{\frac{2c_i}{p_i}}$, and which is the global minimum of f in $\mathbb{R}_+^m$. Thus this case occurs if and only if $\sum_{i=1}^m \ell_i \sqrt{\frac{p_i}{2c_i}} < 1$. In this case, the formula in the statement of the lemma holds for $\lambda^* = 0$.
- Or, the minimum of f in D is reached on the boundary of D. Since $f(\tau) \to \infty$ as soon as some τ_i goes to 0 or to infinity, this minimum has to be on the surface defined by the constraint $g(\tau) = 0$. Then, according to classical minimization techniques, the gradients of f and g are collinear, and moreover they point in the same direction since f increases when τ enters the domain $g(\tau) > 0$. Thus τ^* verifies two equations, for some $\lambda^* \geqslant 0$:

$$(\Sigma)\begin{cases} \mathrm{grad} f(\tau^*) = \lambda^* \, \mathrm{grad} g(\tau^*) \\ g(\tau^*) = 0 \end{cases}$$

  This system can be rewritten by defining $F(\tau, \lambda) = f(\tau) - \lambda g(\tau)$, as follow:

$$(\Sigma) \iff (\tau^*, \lambda^*) \text{ is a critical point of } F$$
$$\iff \begin{cases} (\forall i) \dfrac{\partial F}{\partial \tau_i}(\tau^*, \lambda^*) = 0 \\ \dfrac{\partial F}{\partial \lambda}(\tau^*, \lambda^*) = 0 \end{cases}$$

  We obtain the desired result after a simple algebraic manipulation which shows that F admits a *unique* critical point $\tau^* = \sqrt{\frac{2c_i + \lambda^* \ell_i}{p_i}}$ for some $\lambda^* \geqslant 0$ satisfying $\sum_{i=1}^m \ell_i \sqrt{\frac{p_i}{2c_i + \lambda^* \ell_i}} = 1$.

We have then shown that f admits a unique minimum τ^* on domain D.

[1] Recall that a critical point is a point $\tau^\#$ such that for all i, $\frac{\partial f}{\partial \tau_i}(\tau^\#) = 0$

Now, if all messages have zero broadcast cost, then τ^* verifies:

$$\sum_{i=1}^{m} \frac{\ell_i}{\tau_i^*} = 1 \tag{2}$$

Indeed, if the constraint was not an equality in LB_0, one could decrease one of the τ_i^* and then decrease the objective function in LB_0.

Now, $\tau_i^* = \sqrt{\frac{\ell_i \lambda^*}{p_i}}$ for some $\lambda^* \geqslant 0$. Putting this expression in (2) then gives the expected value of λ^*, from which we compute the values of τ_i^* and LB_0. $\square$

Lemma 6 ($\mathbf{LB_0}$) *The minimization problem* LB_0 *is a lower bound to the cost of any schedule of* $M_1, \ldots, M_m$*. Furthermore, a periodic schedule* S *has cost* LB_0 *if and only if* S *schedules each message* M_i *periodically at intervals of lengths exactly* τ_i^**, where* τ^* *is the unique solution to* LB_0.

Proof. W.l.o.g. (from Lemma 4), we restrict ourselves to periodic schedules. The proof is then along the lines of [AW85, BNBNS98], relaxing the problem by allowing message broadcasts to overlap. Consider a periodic schedule S with period T. We use the results and notations in Corollary 2. The contribution of each message M_i to the cost is: $\mathrm{COST}(S, M_i) = p_i\ell_i + \sum_{j=1}^{n_i} p_i \frac{(t_j^i)^2}{2T} + \frac{n_i}{T} c_i$.

For every given i, since $\sum_{j=1}^{n_i} t_j^i = T$, the sum of the squares is minimized if and only if all terms t_j^i are equal: if $t_j^i = \frac{T}{n_i}$ for all j. Let us define $\tau_i = \frac{T}{n_i}$, we get that:

$$\mathrm{COST}(S, M_i) \geqslant p_i\ell_i + p_i\frac{\tau_i}{2} + \frac{c_i}{\tau_i}$$

Summing up the different contributions yields then a lower bound on the cost of S. Since the total available bandwidth over a period is T, we have: $\sum_{i=1}^{m} n_i\ell_i \leqslant T$, which yields the constraint on the τ_is.

Realizability of the lower bound. According to the proof above, a periodic schedule S has cost $\leqslant \sum_{i=1}^{m} \left(\frac{p_i\tau_i}{2} + \frac{c_i}{\tau_i}\right)$ where $\tau_i = \frac{T}{n_i}$ if and only if it schedules message M_i exactly every τ_i. Since LB_0 has a unique solution τ^* (Lemma 5), a periodic schedule whose cost is equal to LB_0, necessarily broadcasts M_i exactly every τ_i^*. $\square$

Let us say that a lower bound LB for a minimization problem is *tight* if the ratio $\frac{\mathrm{OPT}}{\mathrm{LB}}$ is bounded. The lower bound above was tight in the setting of [BNBNS98]. However, in our model where the message lengths vary, this is no longer the case: $\frac{\mathrm{OPT}}{\mathrm{LB}_0}$ can be arbitrarily large: moreover, using the frequencies $\frac{1}{\tau_i^*}$ as in [AW85, BNBNS98] to broadcast the messages (M_i) can generate arbitrarily bad schedules.

Fact 7 LB_0 *is not tight.*

Proof. Consider the following two messages M_1 and M_2 with zero broadcast cost: M_1 has length 1 and popularity $\frac{L}{L+1}$; and M_2 has length L and popularity $\frac{1}{L+1}$.

Lemma 5 gives: $\mathrm{LB}_0 = \frac{4L}{L+1}$ ($\tau_1^* = 2$ and $\tau_2^* = 2L$). Then $\mathrm{LB}_0 = \Theta(1)$, when L goes to ∞. But according to [BNS99], the optimal schedule alternates scheduling message M_1 $\Theta(L\sqrt{L})$ times and message M_2 once, and has cost $\mathrm{OPT} = \sqrt{L} + O(1) \gg \Theta(1) = \mathrm{LB}_0$.

Furthermore note that if one tries to schedule message M_1 with frequency proportionnal to $\frac{1}{\tau_1^*} = \Theta(1)$ and message M_2 with frequency proportionnal to $\frac{1}{\tau_2^*} = \Theta(\frac{1}{L})$, broadcasting M_2 occupies $L \cdot \Theta(\frac{1}{L}) = \Theta(1)$ of the channel. Then $\Omega(1)$ of the requests for M_1 araise during a broadcast of M_2 and then wait on average $\frac{L}{2}$. The cost of the resulting schedule S is then at least $\frac{L}{2}\Omega(1) = \Omega(L)$ and:

$$\mathrm{COST}(S) = \Omega(L) \quad \gg \quad \mathrm{OPT} = \Theta(\sqrt{L}) \quad \gg \quad \mathrm{LB}_0 = \Theta(1)$$

□

Upon closer examination, one realizes that the cost of a schedule in the previous example is unavoidably high because, whenever a really long message is broadcast, all incoming requests have to wait. This observation led us to the following improved lower bound. We will see in Theorem 16 that this improved lower bound is tight up to a constant factor in the one-channel case.

4.2 A tight lower bound

We present here a tight lower bound LB_α for the data broadcast with non uniform transmission times. Let us define the quantity LB_α as the following minimization problem, for $0 \leqslant \alpha < 1$:

$$\mathrm{LB}_\alpha \begin{cases} \mathbb{E}_p[\ell] \ + \ \min\limits_{\tau > 0} \ \sum\limits_{i=1}^{m} \left((1-\alpha) p_i \dfrac{\tau_i}{2} + \alpha \dfrac{\ell_i^2}{2\tau_i} + \dfrac{c_i}{\tau_i} \right) \\ \qquad \text{subject to: } \sum\limits_{i=1}^{m} \dfrac{\ell_i}{\tau_i} \leqslant 1 \end{cases}$$

Note that for $\alpha = 0$, we obtain LB_0.

As before this minimization problem is solved as follows:

Lemma 8 *For $0 \leqslant \alpha < 1$, the minimization problem LB_α admits a unique solution $\tau^{\circledast}$ verifying:*

$$\tau_i^{\circledast} = \sqrt{\frac{2c_i + \alpha \ell_i^2 + \lambda^{\circledast} \ell_i}{(1-\alpha) p_i}}$$

for a certain $\lambda^{\circledast} \geqslant 0$ (independent of i), such that: $\lambda^{\circledast} = 0$, if $\sum_{i=1}^{m} \ell_i \sqrt{\frac{(1-\alpha)p_i}{2c_i + \alpha \ell_i^2}} \leqslant 1$; otherwise, $\lambda^{\circledast}$ is the unique positive solution to $\sum_{i=1}^{m} \ell_i \sqrt{\frac{(1-\alpha)p_i}{2c_i + \alpha \ell_i^2 + \lambda^{\circledast} \ell_i}} = 1$.

Proof. Straightforward application of Lemma 5 with the modified cost $c_i' = \frac{\alpha \ell_i^2 + 2c_i}{2(1-\alpha)}$. □

Proposition 9 (A tight lower bound) *For any $0 \leqslant \alpha < 1$, the minimization problem LB_α is a lower bound to the cost of any schedule of $M_1, \ldots, M_m$.*

Proof. As before we restrict ourself to periodic schedules. Let S be a periodic schedule. We denote by T its period, and by n_i the number of broadcasts of M_i during a period. We define $\tau_i = \frac{T}{n_i}$. Consider a random request. With probability $\frac{n_i \ell_i}{T} = \frac{\ell_i}{\tau_i}$, it arrives during the broadcast of M_i, and then has to wait at least until the end of the broadcast of M_i, i.e. $\ell_i/2$ on average, plus the length of the message it downloads, i.e. $\mathbb{E}_p[\ell]$ on average. Summing over all messages M_i, we get that:

$$\mathcal{A}\mathrm{ST}(S) \geqslant \sum_{i=1}^{m} \frac{\ell_i}{\tau_i} \cdot \frac{\ell_i}{2} + \mathbb{E}_p[\ell] = \mathbb{E}_p[\ell] + \frac{1}{2} \sum_{i=1}^{m} \frac{\ell_i^2}{\tau_i}$$

A linear $(1-\alpha, \alpha)$ combination of LB_0 with this lower bound on the average service time $\mathcal{A}\mathrm{ST}(S)$ gives us LB_α. □

We will show in section 6 that this lower bound is indeed tight.

5 NP-hardness of the Data-Broadcast problem

We will show that even when there are no broadcast cost, the Data Broadcast problem is already strongly NP-hard.

5.1 Structural properties of an optimal solution

Lemma 10 (Maximum interval) *Let S be a periodic schedule of $M_1, \ldots, M_m$. There exists a periodic schedule S' with* $\mathrm{COST}(S') \leqslant \mathrm{COST}(S)$ *and such that for any i, the intervals between two broadcasts of M_i all have length at most $\mathcal{K}$:*

$$\mathcal{K} = \frac{5\ell_{\max} + 2\sqrt{c_{\max}}}{p_{\min}}$$

Proof. The proof works by induction.

Construction of S_1. Consider in S an interval with maximum length, A_1, between two consecutive broadcasts of the same message. W.l.o.g. this message is M_1 and the first of those two broadcasts of M_1 occurs at time 0. Assume $A_1 \geqslant \mathcal{K}$. We construct a periodic schedule S_1 by inserting in S a broadcast of M_1 at time $\frac{A_1+\ell_1-\ell_{\max}}{2}$, or as soon as possible after that time (that is, if there is a message being broadcast, we wait until the broadcast is finished before inserting M_1).

Analysis of S_1. If T is the period of S, then S_1 has period $T + \ell_1$. We need to show that $\Delta =_{\text{def}} (T + \ell_1)\,\mathrm{COST}(S') - T\,\mathrm{COST}(S) \leqslant 0$.

- Broadcasting M_1 increases the cumulated broadcast cost by c_1.
- In S, the cumulated service time for M_1 has a term $p_1 \frac{A_1^2}{2}$. In S_1, that term disappears and is replaced by:

  $$\frac{p_1}{2}\left(t_1^2 + (A_1 + \ell_1 - t_1)^2\right) \leqslant p_1 \left(\frac{A_1 + \ell_1 + \ell_{\max}}{2}\right)^2$$

  since the insertion time t_1 of M_1 satisfies $\frac{A_1+\ell_1-\ell_{\max}}{2} \leqslant t_1 \leqslant \frac{A_1+\ell_1+\ell_{\max}}{2}$.
- In S, the cumulated service time for M_i ($i \geqslant 2$) has a term $p_i \frac{A_i^2}{2}$, accounting for the interval of length A_i in which M_1 was inserted. In S_1, that term disappears and is replaced by $p_i \frac{(A_i+\ell_1)^2}{2}$.

We obtain:

$$\begin{aligned}
\Delta &\leqslant c_1 + \frac{p_1}{4}\left((A_1 + \ell_1 + \ell_{\max})^2 - 2A_1^2\right) + \sum_{i=2}^{m} \frac{p_i}{2}\left((A_i + \ell_1)^2 - A_i^2\right) \\
&= c_1 + \frac{p_1}{4}\left(-A_1^2 + 2A_1(\ell_1 + \ell_{\max}) + (\ell_1 + \ell_{\max})^2\right) + \sum_{i=2}^{m} \frac{p_i}{2}\left(2\ell_1 A_i + \ell_1^2\right) \\
&\leqslant -\frac{p_1}{4}A_1^2 + \ell_{\max} A_1 + {\ell_{\max}}^2 + c_{\max}
\end{aligned}$$

since $\ell_1 \leqslant \ell_{\max}$, $c_1 \leqslant c_{\max}$, $p_1 + \cdots + p_m \leqslant 1$ and $A_i \leqslant A_1$. In terms of A_1, this last expression is a polynomial of degree 2 with negative leading coefficient. Since $A_1 \geqslant \mathcal{K}$ and $\mathcal{K}$ is larger than the maximum root of that polynomial, we conclude that Δ is negative.

The induction. Repeating the construction, we note that any two insertions of the same message must be at distance at least $(\mathcal{K} - \ell_{\max})/2$ from each other, hence there can be only a finite number of iterations, after which we obtain the requested schedule S'. $\square$

In the rest of this section, we restrict ourselves to the case where all the messages have zero broadcast cost. In that case, it never helps for a schedule to remain idle: one can always decrease the service time of a schedule by skipping all the idle periods. Consequently, we can restrict our search for optimal schedules to schedules that never idle.

The following lemma states that in this situation, there is an optimal schedule which is periodic (note that we currently do not know whether an optimal periodic schedule exists if broadcast costs are present).

Lemma 11 *If the messages have no broadcast costs, then there exists an optimal schedule S^* which is periodic.*

Proof. Since there are no broadcast costs, we only consider schedules with no idle times. We use Lemma 4 again. We will prove that *among periodic schedules*, there exists a schedule S^* of minimum cost. Since:

$$\mathrm{COST}(S^*) = \inf_{S \text{ periodic}} \mathrm{COST}(S) = \inf_{S} \mathrm{COST}(S) = \mathrm{OPT},$$

S^* will be optimal among all schedules. Henceforth we only consider periodic schedules. Moreover, thanks to Lemma 10, we can restrict our search to schedules in which any interval where M_i is not broadcast has length at most $\mathcal{K}$.

We use the penalty-based definition of the cost (see Section 2.3). By Proposition 3:

$$\mathrm{COST}(S) = \mathbb{E}_p[\ell] + \mathcal{A}\mathrm{PE}(S)$$

We use a graph construction similar to [AW87, AGH95]. We consider the following weighted directed graph G with costs. A vertex of G is an m-tuple $\langle a_1, \dots, a_m \rangle$ with $0 \leqslant a_i \leqslant \mathcal{K}$. Semantically, being at that vertex at time t means that for all i, the last broadcast of M_i started at time $t - a_i$. Thus a_i is a sum of values in $\{\ell_1, \ell_2, \dots, \ell_m\}$, hence can only take on a finite number of distinct values; G is thus a finite graph. For every i, there is an edge e from $\langle a_1, \dots, a_i, \dots, a_m \rangle$ to $\langle a_1 + \ell_i, \dots, \ell_i, \dots, a_m + \ell_i \rangle$, of *length* $\ell(e) = \ell_i$. Following that edge means that M_i is the next message broadcast. We associate to that edge a *cost* $c(e)$ which is the sum of the average downloading time $\ell_i \mathbb{E}_p[\ell]$ and of the cumulated penalty over the corresponding period of time. During this period, the average penalty for M_j is: $p_j(a_j + \frac{\ell_i}{2})$, if $j \neq i$, and $p_i \frac{\ell_i}{2}$ if $j = i$. So the cost of e is defined as:

$$c(e) = \ell_i \cdot \left(\mathbb{E}_p[\ell] + \frac{\ell_i}{2} + \sum_{j=1, j \neq i}^{m} p_j a_j \right)$$

The schedules we consider are in bijection with the cycles of G, and the cost of a schedule is exactly the mean cost of the corresponding cycle γ: $\frac{\sum_{e \in \gamma} c(e)}{\sum_{e \in \gamma} \ell(e)}$. Since the cost of every cycle γ of H is the weighted sum of the cost of the elementary cycles composing it, and there are only a finite number of elementary cycles, the minimum is reached by a particular cycle γ^*. The corresponding schedule is optimal. $\square$

5.2 *NP*-hardness

In order to stay in the Turing machine framework for the *NP*-hardness proof, we restrict the input to rational numbers.

Theorem 12 (*NP*-hardness) *The decision problem associated to the restriction of the Data Broadcast problem to the case of messages of rational length and zero cost is strongly NP-hard.*

Proof. We use a reduction of N-partition [GJ79], which, given a sequence of m integers $x_1, \dots, x_m$, decides whether there exists a partition of $\{1, \dots, m\}$ into N sets $I_1, \dots, I_N$ such that: $\sum_{i \in I_1} x_i = \dots = \sum_{i \in I_N} x_i$. Let $\mathcal{S} =_{\text{def}} x_1 + \dots + x_m$. Consider the following instance of Data Broadcast: $(m+1)$ messages $M_1, \dots, M_m, M_{m+1}$ with zero broadcast costs such that:

$$\begin{cases} M_{i \leqslant m}: & \ell_i = x_i \quad \text{and} \quad p_i = \dfrac{x_i}{2\mathcal{S}} \\ M_{m+1}: & \ell_{m+1} = \dfrac{\mathcal{S}}{N^2} \quad \text{and} \quad p_{m+1} = \dfrac{1}{2} \end{cases}$$

Note that all the numbers are rational with polynomial sizes. According to Lemma 5, the solution τ^* of the lower bound LB_0 is:

$$\tau^*_{m+1} = \frac{\mathcal{S}}{N} + \ell_{m+1} \quad \text{and} \quad \tau^*_1 = \dots = \tau^*_m = \mathcal{S} + N\ell_{m+1} = N \cdot \tau^*_{m+1}$$

and the lower bound LB_0 is:

$$\mathrm{LB}_0 = \frac{\sum_{i=1}^m x_i^2}{2\mathcal{S}} + \frac{\mathcal{S}}{4}\left(1 + \frac{1}{N}\right)^2$$

We reduce then N-partition to the following decision problem: (A) *"Does there exist a schedule with average service time less or equal to* $\left(\frac{\sum_{i=1}^m x_i^2}{2\mathcal{S}} + \frac{\mathcal{S}}{4}\left(1 + \frac{1}{N}\right)^2\right)$*?"*.

According to Lemma 11, there exists an optimal periodic schedule of the messages $M_1, \dots, M_m, M_{m+1}$. Thus (A) is equivalent to: (A') *"Does there exist a* periodic *schedule of the messages $M_1, \dots, M_m, M_{m+1}$ with average service time less or equal to* $\left(\frac{\sum_{i=1}^m x_i^2}{2\mathcal{S}} + \frac{\mathcal{S}}{4}\left(1 + \frac{1}{N}\right)^2\right)$*?"*. According to Lemma 6, the answer is "yes" if and only if there exists a periodic schedule S that broadcasts M_{m+1} exactly every $\tau^*_{m+1} = \ell_{m+1} + \frac{\mathcal{S}}{N}$, and each $M_{i \leqslant m}$ exactly every $\tau^*_i = (N\ell_{m+1} + \mathcal{S}) = N \cdot \tau^*_{m+1}$. Such a schedule exists if and only if one can partition $\{M_1, \dots, M_m\}$ into N sets $\mathcal{M}_1, \dots, \mathcal{M}_N$ such that (see Figure 3):

$$\sum_{M_i \in \mathcal{M}_1} \ell_i = \dots = \sum_{M_i \in \mathcal{M}_N} \ell_i$$

□

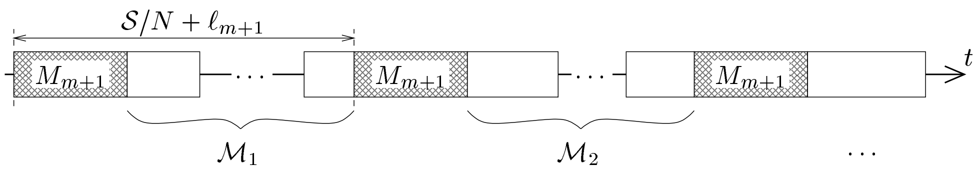

Figure 3: A schedule that satisfies the decision problem (A).

6 A 3-approximation algorithm

6.1 A randomized 3-approximation

According to the lower bound LB_α, one should try to broadcast the messages M_i regularly proportionally to $\tau^{\circledast}_i$. A classical method in designing randomized algorithms (see [Gec83, AW85, BNBNS98, VH97a] in the data broadcasting context) is to choose randomly to broadcast message M_i with probability proportional to $\frac{1}{\tau^{\circledast}_i}$. We will show in this subsection that this leads to a 3-approximation.

Algorithm 1 A randomized scheduler

INPUT:

1. m messages $M_1, \dots, M_m$ with request probabilities $(p_i)_{i=1,\dots,m}$, with lengths $(\ell_i)_{i=1,\dots,m}$ and broadcast costs $(c_i)_{i=1,\dots,m}$.

PREPROCESSING:

Consider $\alpha = 1/3$ and $\tau^{\circledast}$ the unique solution to LB_α:

$$\mathrm{LB}_\alpha \begin{cases} \mathbb{E}_p[\ell] \;+\; \min\limits_{\tau>0} \;\sum\limits_{i=1}^{m}\left((1-\alpha)p_i\dfrac{\tau_i}{2} + \alpha\dfrac{\ell_i^2}{2\tau_i} + \dfrac{c_i}{\tau_i}\right) \\ \qquad\qquad \text{subject to: } \sum\limits_{i=1}^{m} \dfrac{\ell_i}{\tau_i} \leqslant 1 \end{cases}$$

as presented in Lemma 8. Take $\ell_0 = \mathbb{E}_p[\ell]$. Define $\tau_0^{\circledast}$ by $\frac{\ell_0}{\tau_0^{\circledast}} = 1 - \sum_{i=1}^{m} \frac{\ell_i}{\tau_i^{\circledast}}$. For $i = 0, \dots, m$, take:

$$s_i = \frac{\frac{1}{\tau_i^{\circledast}}}{\sum_{j=0}^{m} \frac{1}{\tau_j^{\circledast}}}$$

OUTPUT:

loop

Draw $i \in \{0, \dots, m\}$ with probability s_i.

if $i = 0$ **then**

The server remains idle for an time period of length ℓ_0.

else

The server broadcasts message M_i.

Lemma 13 analyses the expected performance of any algorithm that broadcasts each message M_i with probability s_i and stays idle for a period of time ℓ_0 with probability s_0 (the idle periods are necessary to take broadcast costs into account). Theorem 14 concludes that this algorithm is a 3-approximation when (s_i) and ℓ_0 are chosen as in the preprocessing step.

Lemma 13 *The randomized schedule output by algorithm 1 has expected cost:*

$$\mathbb{E}[\mathrm{COST}(S)] = \frac{\mathbb{E}_s[\ell^2]}{2\;\mathbb{E}_s[\ell]} + \mathbb{E}_p\left[\frac{1}{s}\right] \cdot \mathbb{E}_s[\ell] + \frac{\mathbb{E}_s[c]}{\mathbb{E}_s[\ell]}$$

Proof. We analyse separatly the expected average broadcast cost and the expected average service time. Let $N_i(t)$ the number of broadcasts of message M_i that start during the time period $[0, t]$.

The broadcast cost. By definition, $\mathcal{A}\mathrm{BC}(S, [0, T]) = \sum_{i=1}^{m} \frac{c_i N_i(T)}{T}$. By the Borel-Cantelli theorem, $\lim_{T\to\infty} \frac{N_i(T)}{T} = \frac{s_i}{\mathbb{E}_s[\ell]}$ The linearity of expectation then implies

$$\mathbb{E}[\mathcal{A}\mathrm{BC}(S)] = \sum_{i=0}^{m} \frac{c_i s_i}{\mathbb{E}_s[\ell]} = \frac{\mathbb{E}_s[c]}{\mathbb{E}_s[\ell]}$$

The service time. Consider a request Q for M_i arriving at time t. Its service time $\mathrm{ST}(S, M_i, t)$ is $Y(t) + Z_i + \ell_i$, where $Y(t)$ is the time elapsed from t to the end of the current broadcast at time t, and Z_i is the time elapsed from $t + Y(t)$ to the beginning of the next broadcast of M_i. Remark that $Y(t)$ is independent of i and Z_i is independent of t.

First we compute $\mathbb{E}[Y(t)]$. >From the Borel-Cantelli theorem, the scheduler spends on expectation $\frac{s_j \ell_j}{\mathbb{E}_s[\ell]}$ fraction of the time broadcasting M_j, so with probability $\frac{s_j \ell_j}{\mathbb{E}_s[\ell]}$, request Q

arises during a broadcast of message M_j and then waits $\frac{\ell_j}{2}$ on expectation. Thus:

$$\mathbb{E}[Y(t)] = \sum_{j=0}^{m} \frac{s_j \ell_j}{\mathbb{E}_s[\ell]} \cdot \frac{\ell_j}{2} = \frac{\mathbb{E}_s[\ell^2]}{2\ \mathbb{E}_s[\ell]}$$

Second, to compute $E(Z_i)$, observe that there are on expectation $\sum_{t \geqslant 0} t s_i (1-s_i)^t = \frac{1-s_i}{s_i}$ messages sent before M_i is sent, and each of those messages has expected length $\frac{1}{1-s_i} \sum_{j=0,\ j \neq i}^{m} s_j \ell_j$. So, $\mathbb{E}[Z_i] = \frac{\sum_{j=0,\ j\neq i}^{m} s_j \ell_j}{s_i}$ and a short calculation concludes the proof of the Lemma. □

Theorem 14 *Algorithm 1 is a randomized polynomial* 3*-approximation.*

Proof. Plugging the values of s_i into the Lemma, a straightforward calculation gives the expected cost of the randomized schedule S output by Algorithm 1. Compare this cost to LB_α term by term:

$$\begin{aligned} \mathbb{E}[\mathrm{COST}(S)] &= \sum_{i=1}^{m} p_i \tau_i^{\circledast} + \sum_{i=0}^{m} \frac{\ell_i^2}{2\tau_i^{\circledast}} + \sum_{i=1}^{m} \frac{c_i}{\tau_i^{\circledast}} \\ \mathrm{LB}_\alpha &= \mathbb{E}_p[\ell] + \sum_{i=1}^{m} \left(\frac{1-\alpha}{2}\ p_i \tau_i^{\circledast} + \alpha\ \frac{\ell_i^2}{2\tau_i^{\circledast}} + \frac{c_i}{\tau_i^{\circledast}} \right) \end{aligned}$$

We obtain:

$$\mathbb{E}[\mathrm{COST}(S)] \leqslant \frac{\ell_0}{2} \cdot \frac{\ell_0}{\tau_0^{\circledast}} + \max\left(\frac{2}{1-\alpha}, \frac{1}{\alpha}, 1 \right) \cdot (\mathrm{LB}_\alpha - \mathbb{E}_p[\ell]) \tag{3}$$

Using $\frac{\ell_0}{\tau_0^{\circledast}} \leqslant 1$, $\mathrm{LB}_\alpha \leqslant \mathrm{OPT}$, $\ell_0 = \mathbb{E}_p[\ell]$, and $\alpha = 1/3$, we finally get:

$$\mathbb{E}[\mathrm{COST}(S)] \leqslant 3\,\mathrm{OPT} - \frac{\mathbb{E}_p[\ell]}{2}. \tag{4}$$

□

As a corollary of the proof, we observe that for any $0 < \alpha < 1$, LB_α is a tight lower bound.

6.2 Derandomizing: a deterministic greedy 3-approximation

The algorithm so far constructs a randomized schedule, which has the inconvenience of not being periodic. Periodicity is a particularly useful property when designing cache or prefetching strategy [Ach98] as well as indexing to allow power saving by reducing monitoring time [KZ98]. In this section, we derandomize Algorithm 1 into a deterministic algorithm, which we will then, in the next section, truncate and wrap around cyclically to produce a deterministic periodic schedule with quadratic length.

Our approach is inspired by the analysis in [BNBNS98] of the greedy algorithm of [AW85]. A close look at the potential function used in [BNBNS98] reveals that the algorithm of [AW85] could be obtained from the randomized algorithm of [Gec83] by applying the method of conditional expectations. Similarly, we will apply the method of conditional expectations.

Definition. We denote by $\sigma_i^t(S)$ the time elapsed at time t since the beginning of the last broadcast of M_i before time t in a schedule S. (By convention, we assume that $\sigma_i^t(S) = t$ for all time t before the first broadcast of M_i.) Let $\sigma^t(S) = (\sigma_i^t(S))$ be the *state* of the schedule at time t.

The following lemma shows that the greedy algorithm 2 performs at least as well as the randomized algorithm 1.

Algorithm 2 A greedy scheduler

INPUT: As in Algorithm 1.
PREPROCESSING: As in Algorithm 1.
OUTPUT:

loop

Let σ_i be the time elapsed since the beginning of the last broadcast of M_i.
Choose $i \in \{0, 1, \ldots, m\}$ which minimizes:

$$c_i + \frac{\ell_i^2}{2} - p_i \sigma_i \frac{\mathbb{E}_s[\ell]}{s_i} - \ell_i \underbrace{\left\{ \frac{1}{2}\frac{\mathbb{E}_s[\ell^2]}{\mathbb{E}_s[\ell]} + \frac{\mathbb{E}_s[c]}{\mathbb{E}_s[\ell]} - \sum_{j=1}^{m} p_j \sigma_j \right\}}_{\text{independent of } i}$$

if $i = 0$ **then**
The server remains idle for a time period ℓ_0.
else
The server broadcasts message M_i.

Lemma 15 (Greedy algorithm) *Algorithm 2 outputs a schedule S whose cost is at most the expected cost of Algorithm 1, given in Lemma 13.*

Note 1 *When all the lengths are equal to 1, Algorithm 2 is the algorithm of [BNBNS98].*

Proof. Let M_{g_n} denote the nth message broadcast, where by convention M_0 denotes an idle period of length ℓ_0. We use the penalty-based definition of the cost presented in Section 2.3. According to Equation (1) and Proposition 3:

$$\mathrm{COST}(S) = \widehat{\mathrm{COST}}(S) = \limsup_{N\to\infty} \frac{\widehat{\mathrm{cost}}(1,S) + \cdots + \widehat{\mathrm{cost}}(N,S)}{\ell_{g_1} + \cdots + \ell_{g_N}}, \tag{5}$$

where $\widehat{\mathrm{cost}}(n,S)$ is the sum of $\ell_{g_n}\mathbb{E}_p[\ell]$ and of the penalty and broadcast costs of a schedule S during the broadcast of M_{g_n}.

Let $\sigma^n(S)$ denote the state at the end of the nth broadcast in S. Our analysis uses a potential function defined by:

$$\Phi(n,S) =_{\mathrm{def}} \sum_{j=1}^{m} p_j \sigma_j^n(S) \left(\frac{\mathbb{E}_s[\ell]}{s_j} - \ell_j \right)$$

The definition of Φ generalizes the potential function in [BNBNS98] using the following intuition: $\Phi(n,S)$ is exactly the expected wait of the requests still unserved at the end of the n^{th} broadcast in S, *if* the schedule uses randomized Algorithm 1 after the n^{th} broadcast. A request for M_j would then wait $\sum_{k=0,k\neq j}^{m} \frac{\ell_k s_k}{s_j} = \mathbb{E}_s[\ell]/s_j - \ell_j$ in expectation (see the proof of Lemma 13). Φ thus takes into account the future cost which is not part of the penalty.

Let Avg denote the expected cost of the randomized algorithm, as given in Lemma 13. We will show that the selection rule ensures that at every step n that broadcasts M_i:

$$\widehat{\mathrm{cost}}(n,S) - \Phi(n-1,S) + \Phi(n,S) \leqslant \ell_i \,\mathsf{Avg}, \tag{6}$$

Summing over n and dividing by the sum of the lengths then yields:

$$\frac{\sum_{n=1}^{N} \widehat{\mathrm{cost}}(n,S) + \Phi(N,S) - \Phi(0,S)}{\ell_{g_1} + \cdots + \ell_{g_N}} \leqslant \mathsf{Avg},$$

Algorithm 3 A periodic scheduler

INPUT: As in Algorithms 1 and 2.

PREPROCESSING: As in Algorithms 1 and 2. Let $T_0 = \frac{4\mathcal{L}^2+2\mathcal{C}}{\mathbb{E}_p[\ell]} - \mathcal{L}$.

OUTPUT: The first period of the schedule is as follows.

1. Broadcast each message M_i once ($i = 1 \cdots m$), in order of decreasing s_i.
2. Run Algorithm 2 until time $T_0 + \mathcal{L}$, then finish the current broadcast.

which concludes the proof since $\Phi(N, S)$ is always non-negative and $\Phi(0, S) = 0$.

We now concentrate on proving equation (6). Let S_i be the schedule which is identical to S for the first $(n-1)$ steps but broadcasts message M_i during the n^{th} step. We will show two facts: first, that equation (6) holds on average when i is selected according to distribution s; and second, that the greedy selection rule selects the i that minimizes the difference between the left and the right hand sides in equation (6), thus ensuring that (6) holds.

In S_i, we have $\sigma_j^n(S_i) = \sigma_j^{n-1}(S) + \ell_i$ for $j \neq i$, and $\sigma_i^n(S_i) = \ell_i$. A simple calculation gives then:

$$\Phi(n, S_i) - \Phi(n-1, S_i) = \ell_i \sum_{j=1}^{m} p_j \left(\frac{\mathbb{E}_s[\ell]}{s_j} - \ell_j \right) - p_i \sigma_i^{n-1}(S) \left(\frac{\mathbb{E}_s[\ell]}{s_i} - \ell_i \right)$$

Moreover, the average penalty for M_j during the broadcast of M_i is $p_j(\sigma_j^{n-1}(S) + \frac{\ell_i}{2})$, if $j \neq i$, and $p_i \frac{\ell_i}{2}$, if $j = i$. Thus:

$$\widehat{\text{cost}}(n, S_i) = c_i + \frac{\ell_i^2}{2} + \ell_i \sum_{j=1, j\neq i}^{m} p_j \sigma_j^{n-1}(S) + \ell_i \mathbb{E}_p[\ell]$$

Thus, the left-hand side of equation (6) equals:

$$L.H.S.(6) = c_i + \frac{\ell_i^2}{2} + \ell_i \sum_{j=1}^{m} p_j \sigma_j^{n-1}(S) + \ell_i \mathbb{E}_p\left[\frac{1}{s}\right] \cdot \mathbb{E}_s[\ell] - p_i \sigma_i^{n-1}(S) \frac{\mathbb{E}_s[\ell]}{s_i} \tag{7}$$

Taking the expectation over i gives:

$$\mathbb{E}_s[L.H.S.(6)] = \mathbb{E}_s[c] + \frac{1}{2}\mathbb{E}_s[\ell^2] + \mathbb{E}_p\left[\frac{1}{s}\right] \cdot \mathbb{E}_s[\ell]^2 = \mathbb{E}_s[\ell]\,\mathsf{Avg},$$

which concludes the first fact. Second, using equation (7) and the expression of Avg given in Lemma 13, note that the difference between the left and the right hand sides in equation (6) can be rewritten as

$$-p_i \sigma_i^{n-1}(S) \frac{\mathbb{E}_s[\ell]}{s_i} + c_i + \frac{\ell_i^2}{2} - \ell_i \left\{ \frac{\mathbb{E}_s[\ell^2]}{2\mathbb{E}_s[\ell]} + \frac{\mathbb{E}_s[c]}{\mathbb{E}_s[\ell]} - \sum_{j=1}^{m} p_j \sigma_j^{n-1}(S) \right\}$$

The greedy rule then selects at each step the message M_i that minimizes this quantity. □

6.3 Truncating: a deterministic periodic 3-approximation

In this section, we prove that one can stop the greedy algorithm after a quadratic number of steps and obtain a periodic 3-approximation by wrapping the schedule around.

The construction is illustrated on Figure 4. Algorithm 3 defines a periodic schedule of period $T + \mathcal{L}$, with $T_0 \leqslant T < T_0 + \ell_{\max}$.

Theorem 16 (Deterministic and periodic approximation) *There exists a deterministic polynomial algorithm that outputs a periodic schedule S with period $\leqslant \frac{4\mathcal{L}^2+2\mathcal{C}}{\mathbb{E}_p[\ell]}$ and with cost $\leqslant 3\,\mathrm{OPT}$.*

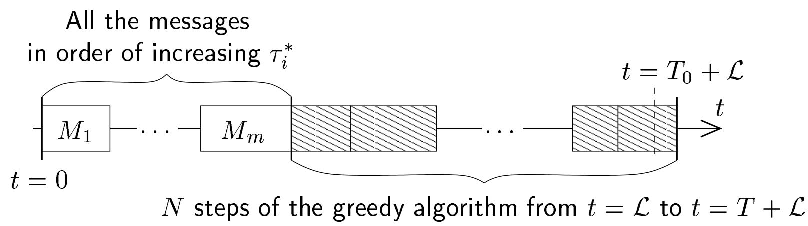

Figure 4: The periodic approximation.

Proof. We adopt the same notations ($\sigma^n(S)$, $\Phi(n, S)$,...) as in the proof of Lemma 15. As in that Lemma, we use the penalty-based definition of the cost of S (see Section 2.3). S has period $(T + \mathcal{L})$.

Analysis of the time interval $[\mathcal{L}, T + \mathcal{L}]$. According to Equation (6), the cumulated cost of S for this time interval is bounded by:

$$T\,\mathsf{Avg} + \Phi(m, S) - \Phi(N + m, S), \tag{8}$$

where N is the number of broadcasts of S during that time interval.

Analysis of the time interval $[0, \mathcal{L}]$. Every message is broadcast once during this period, the total broadcast cost for this period is then clearly $\mathcal{C}$. If t_i denotes the time of the first broadcast of M_i, then the cumulated penalty of S during $[0, \mathcal{L}]$ is:

$$\sum_{i=1}^{m} p_i \left(t_i \sigma_i^{N+m}(S) + \frac{t_i^2}{2} + \frac{(\mathcal{L} - t_i)^2}{2} \right) \leqslant \sum_{i=1}^{m} p_i t_i \sigma_i^{N+m}(S) + \mathcal{L}^2.$$

Since the messages are broadcast by order of decreasing s_i, t_i is less than the expected time of the first broadcast of M_i in Algorithm 1, and so $\sum_i p_i t_i \sigma_i^{N+m}(S) \leqslant \Phi(N + m, S)$.

Adding the cumulated costs during the two time intervals and using Equation (5), we obtain that the total cost satisfies:

$$(T + \mathcal{L})\,\mathrm{COST}(S) \leqslant T\,\mathsf{Avg} + \Phi(m, S) + \mathcal{L}^2 + \mathcal{C}$$

Now, $\Phi(m, S)$ is the expected future waiting time of requests still unserved at the beginning of the period, if one was to continue with the randomized algorithm. All the requests arose in the past $\mathcal{L}$ time, so $\Phi(m, S) \leqslant \mathcal{L}\,\mathsf{Avg}$. Thus:

$$\mathrm{COST}(S) \leqslant \mathsf{Avg} + \frac{\mathcal{L}^2 + \mathcal{C}}{T + \mathcal{L}}$$

From our definition of T_0, $T + \mathcal{L} \geqslant T_0 + \mathcal{L} \geqslant \frac{4\mathcal{L}^2 + 2\mathcal{C}}{\mathbb{E}_p[\ell]}$. Recalling from Equation (4) that $\mathsf{Avg} \leqslant 3\,\mathrm{OPT} - \frac{\mathbb{E}_p[\ell]}{2}$, we conclude that:

$$\mathrm{COST}(S) \leqslant 3\,\mathrm{OPT}.$$

Analysis of processing time. The computation of τ^* will be discussed in the next section. The time complexity here is the number of steps of the greedy algorithm N. Time increases at each step by at least by $\ell_{\min}$, so the schedule is computed after $N \leqslant \frac{4\mathcal{L}^2 + 2\mathcal{C}}{\mathbb{E}_p[\ell]\ell_{\min}}$ steps. □

6.4 Precision issues: Computing an approximation of $\tau^{\circledast}$

Solving the minimization problem LB_α reduces to find the unique positive solution $\lambda^{\circledast}$ to:

$$\sum_{i=1}^{m} \ell_i \sqrt{\frac{a_i}{b_i + \lambda \ell_i}} = 1$$

where $\ell_i > 0$, $0 < a_i \leqslant 1$ and $b_i > 0$. Since $\lambda \mapsto \sum_{i=1}^{m} \sqrt{\frac{a_i}{b_i+\lambda}}$ is a convex decreasing function, Newton algorithm is a very efficient solution. We have $1 \leqslant \sum_{i=1}^{m} \sqrt{\frac{a_i \ell_i}{\lambda^{\circledast} + \min_i \frac{b_i}{\ell_i}}}$. Since $a_i \leqslant 1$ and $b_i \geqslant 0$, we get:

$$\lambda^{\circledast} \leqslant \left(\sum_{i=1}^{m} \sqrt{\ell_i} \right)^2 \leqslant m\mathcal{L}$$

Thus Newton algorithm computes an approximation of $\lambda^{\circledast}$ within ε accuracy from the seed $\lambda = 0$ in time $O(\log\log\{\frac{1}{\varepsilon}, \mathcal{L}, m\})$. We can thus compute in sublogarithmic time an approximation of $\tau^{\circledast}$ within ε accuracy from which the algorithm in Theorem 16 computes a periodic schedule with cost $\leqslant 3\,\mathrm{OPT} + O(\varepsilon)$.

7 A multiple channels framework

The problem as stated since the beginning of this paper can be seen a *single channel* data broadcast: the server can only broadcast one type of message at a time. One can naturally extend similarly to [BNBNS98] the definitions to a multiple channels framework where the scheduler has W channels available for broadcasting the messages:

- *A schedule* S of $M_1, \ldots, M_m$ on W channels is, formally, a set of ordered pairs $(s_w(n), t_w(n))$, where $1 \leqslant w \leqslant W$. For any $n \geqslant 1$, we define $s_w(n) = i$ if M_i is the nth message broadcast on channel w, and the starting time of this broadcast is $t_w(n)$. Since channel w must finish broadcasting a message before it can start broadcasting any other message, we have:

  $$(\forall w)\ (\forall n) \quad t_w(n+1) \geqslant t_w(n) + \ell_{s_w(n)}$$

  S is *periodic* if there is a $T > 0$ (the period) such that for every channel w, there exists $N_w > 0$ such that for any $n \geqslant 1$, $s_w(n+N) = s_w(n)$ and $t_w(n+N) = t_w(n) + T$.

- *The clients* connect at a uniform random time (given by some Poisson process), ask for a random message M_i (according to the distribution p_i), monitor all the channels and download the requested message as soon as it is broadcast on some channel (clients cannot start downloading in the middle of the broadcast of a message but have to wait until the next broadcast).

- *The cost* is defined as before as the sum of the average service time (average monitoring time plus average downloading time) and of the average broadcast cost.

This case is widely not understood. If finding an optimal periodic schedule is clearly NP-hard, no approximation algorithm with performance guarantee are known at this point. The main point is that the lower bound techniques used in Section 4.2 do not work. The following sections proposes a tour by examples of some of the difficulties we have noted. As we will see the multi-channels non-uniform transmission time case differs radically from the multi-channel uniform case in [BNBNS98].

7.1 What is still true?

Expression and reversibility of the cost. The expression of the cost given in Lemma 1 simply rely on the time elapsed two consecutive broadcasts of the same messages. Hence, the cost has the same expression in the multiple channel framework (beware that the two consecutive broadcasts may not occur on the same channel) as can be seen on Figure 5. The reversibility lemma (Proposition 3) holds for the same reason as well in the multiple channel framework.

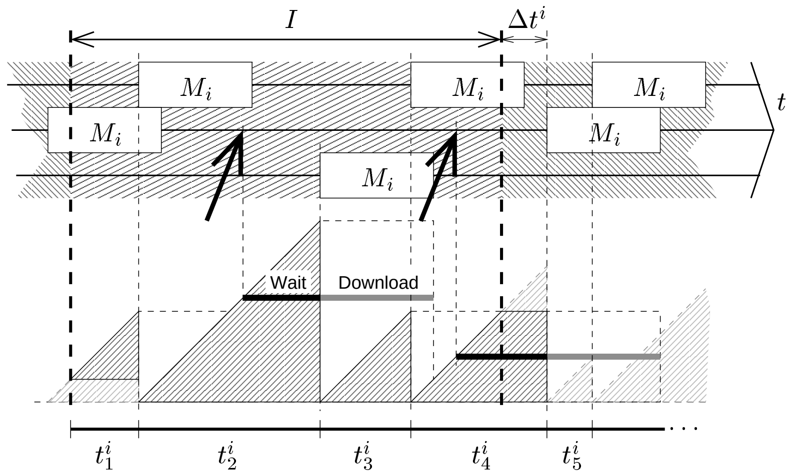

Figure 5: The cost has the same represention in the multiple channel framework.

LB_0 **is not tight.** The lower bound LB_0 extends naturally to a lower bound on the cost of every schedule on W channels by increasing the available bandwidth from 1 to W. We get:

$$\mathrm{LB}_0 \begin{cases} \mathbb{E}_p[\ell] \; + \; \min\limits_{\tau>0} \; \sum\limits_{i=1}^{m} \left(p_i \dfrac{\tau_i}{2} + \dfrac{c_i}{\tau_i} \right) \\ \text{subject to: } \sum\limits_{i} \dfrac{\ell_i}{\tau_i} \leqslant W \end{cases}$$

Again, this lower bound is not tight. We extend indeed the example in Fact 7 as follow: consider $W+1$ messages $M_1, \ldots, M_{W+1}$ where M_i has length $\ell_i = L^{i-1}$ and popularity $p_i = \alpha \cdot L^{1-i}$ where α is chosen so that $p_1 + \cdots + p_{W+1} = 1$.

Fact 17 *For the set of messages above, the optimal schedule on W channels has cost* $\mathrm{OPT} = \Theta\big(L^{2^{-W}}\big)$ *but* $\mathrm{LB}_0 = \Theta(1)$ *when L goes to infinity.*

Proof sketch. A simple adaptation of lemma 5 gives that $\mathrm{LB}_0 = \Theta(1)$. But the optimal schedule on W channels has cost $\Theta\big(L^{2^{-W}}\big) \gg \mathrm{LB}_0$. It is a periodic schedule with period $L^{W+2^{-W}}$ and constructed recursively as follows: between time $t = 0$ and $t = L^W$, M_{W+1} is broadcast on channel W and $M_1, \ldots, M_W$ are broadcast on the $W-1$ first channels according to the optimal schedule on $W-1$ channels; and between time $t = L^W$ and $t = L^{W+2^{-W}}$, messages $M_1, \ldots, M_W$ are broadcast continously each on one of the W channels. The complete proof can be found in [Sch00a]. □

NP-**hardness.** The reduction from N-partition given in Theorem 12 extends simply by adding $W-1$ messages with popularities and lengths adjusted so that each would be broadcast continously on one of the channels $2, \ldots, W$ according to LB_0 (see [Sch00a]). This gives that finding an optimal cost of a periodic schedule on W channels is strongly NP-hard.

7.2 What is not true anymore?

lower bounds. LB_α ($alpha > 0$) is not a lower bound for the multiple channel case: the argument that no request can be served during the broadcast of a message does not hold. Unlike in the uniform setting in [BNBNS98], the optimal cost on W channels depends non-trivialy on W. As can be seen with the example of Fact 17, adding a channel can reduce the optimal cost by an arbitrary factor: on W channels the optimal cost for the $W+1$ messages is $\Theta(L^{2^{-W}})$ which tends to infinity when L grows; but with $W+1$ channels, the schedule that broadcasts M_i constinously on channel i has cost $\Theta(1)$. In the uniform setting only a factor of $\frac{W}{W+1}$ could be saved. As far as we know, no reliable lower bounds are known for the multiple channel setting.

Periodicity of an optimal schedule. Unlike the uniform setting, periodicity of an optimal schedule cannot be guaranteed when broadcasting on several channels. Consider the scheduling problem on two channels of two messages M_1 and M_2, of lengths $\ell_1 = 1$ and $\ell_2 = \sqrt{2}$, and request probabilities $p_1 = \frac{\sqrt{2}}{1+\sqrt{2}}$ and $p_2 = \frac{1}{1+\sqrt{2}}$, and zero costs. A direct modification of Lemma 6 (Section 4.1) gives that $\mathrm{LB}_0(M_1, M_2) = \frac{\sqrt{2}}{2(1+\sqrt{2})}$ is a lower bound on the cost of any twp-channels schedule of M_1 and M_2. Note that it is realized by the schedule that never stops broadcasting M_1 on the first channel and M_2 on the second.
But, if there was a periodic optimal schedule, it should, according to Lemma 6, broadcast M_1 every $\tau_1^* = 1$ time and M_2 every $\tau_2^* = \sqrt{2}$ time: impossible since $\frac{1}{\sqrt{2}}$ is irrational. We don't whether having only rationnal lengths implies the existence of an optimal periodic schedule, but this condition was not required in the single channel case.

Unlike the uniform setting, when broadcasting on several channels, we cannot assume that an optimal schedule never idles even if the message have zero broadcast cost (as in Sectionsec:optimal:periodic). Consider the broadcast of following two messages M_1 and M_2 on two channels: $\ell_1 = 2$ and $p_1 = 1 - \varepsilon$; $\ell_2 = 3$ and $p_2 = \varepsilon$, for some parameter ε. For ε small enough, M_2 will be broadcast rarely. M_1 is the important message. Between two consecutive broadcasts of M_2, M_1 has to be broadcast according to an optimal pattern, that is to say: all the time, on both channels with an offset of one unit of time between the channels (see Figure 6). Suppose that we are given a schedule that never idles: after each broadcast of M_2 on some channel, either the pattern for M_1 is broken inducing a huge increase of the cost, or M_2 is broadcast twice in a row delaying the resuming of the pattern which increases the cost as well. In fact, none of the schedules which never idles is optimal. But an optimal schedule idles for one unit of time around every broadcast of M_2 as shown on Figure 6.

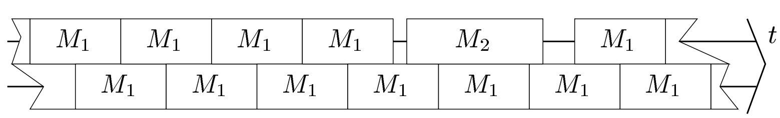

Figure 6: An optimal schedule.

The condition of the existence of an optimal solution in the multiple setting is open. As a consequence, it is not even clear *a priori* whether an optimal solution can be described in finite time.

8 Conclusion

As far as we know this work is the first yielding a constant factor approximation for the non-uniform length case (on one channel). The multiple channel case remains still largely not understood. Finding a good (tight) lower bound for the multiple channel case with non-uniform length messages is certainly an interesting challenging question. A preemptive model for the non-uniform length broadcast problem is also proposed in [Sch00b]. It turns out that the problem with preemption is much better understood. In particular, it shows that if preemption is allowed, one can reduce the expected waiting time by an arbitrary factor.

This paper and others focused on minimizing the expected waiting time, however one might also want to provide some guarantee in terms of quality of service. For example, an interesting extension of the problem would like consist in ensuring that the maximum waiting time is never larger than some bound at least for some "important" messages. An other possible extension would try to schedule at the same time push requests for the most popular messages (based on popularity distribution) and pull requests for less popular messages (based on the real request of the users for these messages); as far as we know this problem is currently unexplored.

References

[AAFZ95] S. Acharya, R. Alonso, M. J. Franklin, and S. Zdonik. Broadcast disks: data management for asymetric communications environment. In *Proc. of ACM International Conference on Management of Data (SIGMOD'95)*, Santa Cruz, CA, June 1995. Available at www.cs.brown.edu/research/ddg.

[Ach98] Swarup Acharya. *Broadcast Disks: Dissemination-based Management for Assymmetric Communication Environments*. PhD thesis, Brown University, May 1998. Available at www.bell-labs.com/~acharya.

[AFZ97] A. Acharya, M. Franklin, and S. Zdonik. Balancing push and pull for data broadcast. In *Proc. of ACM International Conference on Management of Data (SIGMOD97)*, Tuscon, Arizona, Mai 1997. Disponible à www.cs.brown.edu/research/ddg.

[AGH95] S. Anily, C. A. Glass, and R. Hassin. The scheduling of maintenance service. 1995. unpublished.

[Amm87] Mostafa H. Ammar. Response time in a teletext system: An individual user's perspective. *IEEE Transactions on Communications*, COM-35,11:1159–1170, Nov. 1987.

[AW85] M. H. Ammar and J. W. Wong. The design of teletext broadcast cycles. In *Performance Evaluation*, volume 5(4), pages 235–242, 1985.

[AW87] M. H. Ammar and J. W. Wong. On the optimality of cyclic transmission in teletext systems. In *IEEE Trans. on Comm.*, volume COM-35(11), pages 1159–1170, 1987.

[Bes96] Azer Bestavros. Speculative data dissemination and service to reduce server load, network traffic and service time in distributed information systems. In *Proc. of the 1996 International Conf. on Data Engineering (ICDE'96)*, March 1996. Available at cs-www.bu.edu/~best.

[BNBNS98] A. Bar-Noy, R. Bhatia, J. (S.) Naor, and B. Schieber. Minimizing service and operation costs of periodic scheduling. In *Proc. of the 9th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA'98)*, pages 11–20, 1998.

[BNS99] A. Bar-Noy and Y. Shilo. Optimal broadcasting of two files over an asymetric channel. In *Proc. of INFOCOM'99*, 1999.

[Gec83] J. Gecsei. The architecture of videotex systems. *Prentice Hall, Englewood Cliffs, N.J.*, 1983.

[GJ79] M. R. Garey and D. S. Johnson. *Compters and intractability, A guide to the theory of NP-completeness.* W. H. Freeman and Co., 1979.

[KS99] C. Kenyon and N. Schabanel. The data broadcast problem with non uniform time. In *Proc. of the 10th Symp. on Discrete Algorithms (SODA'99)*, pages 547–556, Jan. 1999. Available at www.ens-lyon.fr/∼nschaban.

[KSY00] C. Kenyon, N. Schabanel, and N. Young. A polynomial-time approximation scheme for the data broadcast problem. In *Proc. of the 32nd ACM Symp. on Theory of Computing (STOC 2000)*, pages 659–666, Portland, May 2000. Available at www.ens-lyon.fr/∼nschaban.

[KZ98] S. Khanna and S. Zhou. On indexed data broadcast. In *Proceedings of 30th STOC*, volume 30, pages 463–472, 1998.

[Sch00a] N. Schabanel. Algorithmes d'approximation pour les télécommunications sans fils: Ordonnancement pour la dissémination de données et allocation statique de fréquences. PhD Thesis n°00ENSL0145, LIP, ENS-Lyon, France, Jan 2000. Available (in French) at www.ens-lyon.fr/∼nschaban.

[Sch00b] N. Schabanel. The data broadcast problem with preemption. In *LNCS 1770 Proc of the 17th Int. Symp. on Theoretical Computer Science (STACS 2000)*, pages 181–192, Feb. 2000. Available at www.ens-lyon.fr/∼nschaban.

[ST97] C.-J. Su and L. Tassiulas. Broadcast scheduling for information distribution. In *Proc. of INFOCOM*, 1997.

[VH97a] N. Vaidya and S. Hameed. Efficient algorithms for scheduling single and multiple channel data broadcast. Tech. Report 97-002, Texas A&M univ., Feb. 1997.

[VH97b] N. Vaidya and S. Hameed. Log time algorithms for scheduling single and multiple channel data broadcast. In *Proc. of the 3rd ACM/IEEE Conf. on Mobile Computing and Networking (MOBICOM)*, Sep. 1997.