The Data Broadcast Problem with Non-Uniform Transmission Times

Claire Kenyon

LRI, Université Paris XI
Bâtiment 490
F-91405 Orsay, France
Claire.Kenyon@lri.fr
http://www.lri.fr/∼kenyon

Nicolas Schabanel

LIP, ENS Lyon
46, allée d'Italie
F-69364 Lyon Cedex 07, France
Nicolas.Schabanel@ens-lyon.fr
http://www.ens-lyon.fr/∼nschaban

Abstract. The data broadcast problem consists in finding an infinite schedule to broadcast a given set of messages so as to minimize the average response 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. The structural properties of the optimal solutions appear to be very different from the uniform case. We prove that the data broadcast problem is *NP*-hard, even if the broadcast costs are all zero, and give a randomized 3-approximation algorithm for broadcasting messages on a single channel.

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 happens when mobile clients retrieve information from a server basestation through a wireless medium. In broadcasting protocols, items are scheduled over the infinite time-horizon. The requests do not propagate in the system, but wait for the requested items to be broadcast, thus making the system *pseudointeractive* : the schedule is independent of the incoming requests, since it is oblivious to them. An efficient broadcast scheme seeks to minimize the average response 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 modelled similarly.

1.2 Problem definition

In this paper we will adopt the Data-broadcast terminology. Given C broadcast channels, the input consists of m messages M_i, $i = 1..m$, 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 the infinite time-horizon, so as to minimize the sum of the average response time to Poisson requests and of the average broadcast cost, i.e. so as to minimize $\limsup_{T\to\infty} \mathrm{ART}(S,[0,T]) + \mathrm{BC}(S,[0,T])$; here $\mathrm{ART}(S,[0,T])$ denotes the average response time to a request which is generated at a random uniform instant between 0 and T, requests message M_i with probability p_i, and must then wait until the start of the next broadcast of M_i; and $\mathrm{BC}(S,[0,T])$ 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 litterature on data broadcasting (see Lemma 1). In this paper, we mostly treat the single channel case ($C = 1$).

In the Maintenance Scheduling problem, given C independent operators, the input consists of m machines M_i, $i = 1..m$, defined by the *length of their maintenance* ℓ_i, 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 the 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 Lemma 1, identical to the Data Broadcast problem.

The Multi-item Replenishment problem is a special case of the maintenance scheduling problem [4].

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 equal to 1.

Let us first review results on the single-channel discrete-time and uniform-lengths Data Broadcast problem without broadcast costs. This problem was first studied in the context of Teletext. In [8], 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 [2], Ammar and Wong study periodic broadcast schedules. They derive an algebraic expression for the average response time (which is essentially Lemma 1 below), and prove a lower bound (essentially the one in Lemma 5 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 [3], 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 and is based on [12] and [15, pp.510–511]. (They do not provide a performance ratio analysis). Further work [1] attacks a model where successive requests are not independant and prefetching may be used. Greedy heuristics for the same problem are empirically tested in [16, 4]. This study on the single-channel uniform-length case, was pursued by [4], who obtained a deterministic 2.5-approximation algorithm, and moreover by [5] who obtained a deterministic 2-approximation algorithm and proved that the golden ratio heuristic of [3] has approximation ratio 9/8. Finally in [14], the authors study indexed data broadcast where the objective function to minimize is a combination of response time and tuning time.

Even in the uniform length model, very few articles consider the multi-channel problem. Previous work for the general uniform-length Maintenance Scheduling problem can be found in [5]. 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 $O(1)$-approximation algorithms for the Data Broadcast and the Maintenance Scheduling problems, even when there is more than one channel.

[17, 18] are the only references which considers non-uniform length messages: they report some experimental results for heuristics on one or two channels.

Related work can also be found in [6, 7, 11, 10, 9, 13].

1.4 The results

The main results of this paper concern the one-channel case when the message lengths (transmission times) are not all equal. This generalization is significantly harder: in [3], 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 Example in Section 4.2); furthermore as opposed to the uniform length case studied in [5], the ratio of the costs of two optimal schedules of the same set of messages on C and $C+1$ channels is no more bounded by $(C+1)/C$ but can be arbitrarily large (See example page 6). In general it is not even clear *a priori* whether the optimal solution exists or 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".

First, we prove that when there are no broadcast costs, there exists an optimal schedule which is periodic; however, we then show that, even in that simpler setting, the problem is *NP*-hard (Note that in [5] *NP*-hardness is proved only when broadcast costs are present.) One delicate point in the *NP*-hardness proof is that the input parameters are real numbers. Thus, in order to stay with the standard Turing machine computation model, we prove *NP*-hardness assuming that the parameters are rational numbers.

Then, we present a randomized polynomial approximation for the problem whose asymptotic performance ratio is 3, and whose absolute performance ratio is $3+\varepsilon$ (where ε can be made arbitrarily small). To design that algorithm, we first observe that, as stated by the authors of [5], the natural extension of the lower bound of [2] to our model is no longer tight in the non-uniform-length case, and can, in fact, be arbitrarily bad. One important ingredient of our algorithm here consists in designing a second lower bound which, when suitably combined with the one derived from [2], yields a lower bound which is tight up to a constant factor. This is the key to finding a constant-factor approximation.

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. In Section 3, we prove structural properties of optimal solutions which will be useful for the following. Section 4 concentrates on the one-channel case which is proved to be *NP*-complete even if zero broadcast cost are assumed. It also shows that there exists a periodic optimal solution, if zero broadcast costs are assumed, and gives a finite-time algorithm for constructing it. We finally give in Section 5 a 3-approximation algorithm for the one-chanel case.

We are optimistic on the possibility of designing a deterministic algorithm and are currently working on extending our work to the multi-channel setting.

2 Definitions and Notations

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

The schedule. A *schedule* S of $M_1, \dots, M_m$ on C channels is, formally, a set of ordered pairs $(s_c(n), t_c(n))$, where $1 \leqslant c \leqslant C$. For any $n \geqslant 1$, we define $s_c(n) = i$ if M_i is the nth message broadcast on channel c, and the starting date of this broadcast is $t_c(n)$. Since channel c must finish broadcasting a message before it can start broadcasting any other message, we have:

$$(\forall c)\ (\forall n) \quad t_c(n+1) \geqslant t_c(n) + \ell_{s_c(n)}$$

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

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

The second quantity is the *broadcast cost* $\mathrm{BC}(S)$ 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 broadcast cost of S over I, $\mathrm{BC}(S, I)$, is defined as the sum of the cost of all the messages whose broadcast begins in I, divided by the length of I. $\mathrm{BC}(S)$ is then defined as:

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

Remark that one could have chosen another definition for $\mathrm{BC}(S, I)$ by taking into account the messages whose broadcast starts in I but ends outside I in a different way ; but this is unimportant since it does not affect $\mathrm{BC}(S)$. The quantity which we want to minimize is the *cost* of a schedule S which we define as follows:

$$\mathrm{COST}(S) = \mathrm{ART}(S) + \mathrm{BC}(S)$$

Note that up to scaling the costs c_i, any linear combinaison of ART and BC can be considered.

Lemma 1 introduces notations which will be used throughout the paper and relates our definition of cost to the algebraic definition given in previous work [2, 5] 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 response time to the request is:*

$$\mathrm{ART}(S, I) = \frac{1}{2|I|} \sum_{i=1}^{m} p_i \left\{ \sum_{j=1}^{n_i+1} (t_j^i)^2 - (\Delta t^i)^2 \right\}$$

and the broadcast cost of S over I is:

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

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 $t_{j \geqslant 2}^i$ 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. The average response time to a request for M_i in S is given by $\frac{1}{|I|} \int_I \mathrm{wait}_i(t) dt$, where $\mathrm{wait}_i(t)$ is the interval from instant t to the next broadcast of M_i. But for $t \in [t_{j-1}^i, t_j^i]$, $\mathrm{wait}_i(t) = t_j^i - t$. Thus $\mathrm{wait}_i(t)$ is piecewise affine and a simple calculation leads to the result. □

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

$$\mathrm{COST}(S) = \frac{1}{T} \sum_{i=1}^{m} n_i c_i + \frac{1}{2T} \sum_{i=1}^{m} p_i \sum_{j=1}^{n_i} (t_j^i)^2$$

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.

3 Technical results

3.1 Reductions

Previous work relied on the existence of a periodic optimal schedule. In our setting, this is no longer true (with the exception of Lemma 8), and could in principle make all the manipulations of infinite schedules quite tricky. The way around this problem is given by Lemma 3, below: periodic schedules approach optimal schedules to arbitrary accuracy. This kind of *density property* will be used to focus on periodic schedules, as it implies, first, that lower bounds on the cost of periodic schedules are also lower bounds on any schedule (see Lemma 5 and 6); second, that the existence of an optimal schedule among periodic schedules implies the existence of an optimal schedule, which is periodic (see Lemma 8).

Lemma 3 (Density of periodic schedules) *Let S be a schedule of $M_1, \ldots, M_m$ on C channels. For any $\varepsilon > 0$, there exists a* periodic *schedule S' whose cost satisfies:*

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

Proof sketch. Deferred to the Appendix page 11. □

The next lemma shows that we can restrict ourselves to schedules where every message is broadcast fairly frequently.

Lemma 4 (Maximum interval) *Let S be a periodic schedule of $M_1, \ldots, M_m$ on C channels. There exists a periodic schedule S' such that for any i, any interval where M_i is not broadcast has length at most K_i, where:*

$$K_i = \frac{3\mathcal{L}}{p_i} + 4\mathcal{L} + \frac{2\ \mathcal{C}}{\mathcal{L}}, \text{ with } \mathcal{L} = \sum_{j=1}^{m} \ell_j,\ \mathcal{C} = \sum_{j=1}^{m} c_j$$

and such that: $\mathrm{COST}(S') \leqslant \mathrm{COST}(S)$

Proof sketch. Construction of S'. We denote by T the period of S. By induction on i we just need to show the result for M_1. Let us consider 2 consecutive occurrences of M_1 at distance greater than K_1, and let $K = K_1/2$. W.l.o.g., the first broadcast of M_1 starts at time $t = 0$. The idea is to insert a broadcast of M_1 on some channel at time K. This has to be done a little bit carefully: first, every broadcast on that channel after time K will be staggered by ℓ_1, so, to keep the channels properly synchronized, we need to insert a gap of length ℓ_1 in all the other channels at time K. Second, we cannot do this at exactly time K, since the channels may be in the middle of

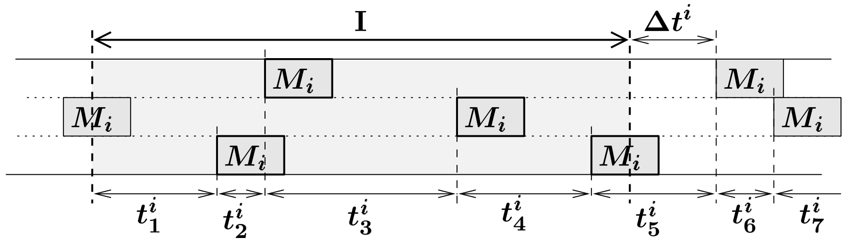

Figure 1: Illustration of the notations used in Lemma 1. Here, $n_i = 4$.

broadcasting something, so we have to wait; third, to avoid delaying the other messages unduly, we also broadcast them at (or around) time K, in order of first previous appearance. The construction of S' is illustrated on Figure 2. Its detailed construction and analysis are somewhat technical and deferred to the Appendix page 11. □

3.2 Lower bounds

Finding good lower bounds is a key point to designing provably efficient algorithms for this problem. The first lower bound below is a straightforward generalization of the one used in [5].

Lemma 5 ($\mathbf{LB_0}$) *The following minimization problem is a lower bound to the cost of a schedule of $M_1, \ldots, M_m$ on C channels:*

$$\mathrm{LB}_0 \left\{ \begin{array}{l} \min\limits_{\tau > 0} \dfrac{1}{2} \sum\limits_{1 \leqslant i \leqslant m} p_i \tau_i + \dfrac{2c_i}{\tau_i} \\ \text{subject to} \sum\limits_i \dfrac{\ell_i}{\tau_i} \leqslant C \end{array} \right.$$

This minimization problem admits a unique solution τ^* verifying: $\tau_i^* = \sqrt{\frac{2c_i + \lambda^* \ell_i}{p_i}}$ for a certain $\lambda^* \geqslant 0$. Furthermore, this lower bound is realized by a periodic schedule S if and only if M_i is scheduled in S periodically at intervals of lengths exactly τ_i^*.

Proof. From Lemma 3, periodic schedules are dense in the set of all schedules, therefore

$$\inf_{S \text{ periodic}} \mathrm{COST}(S) = \inf_S \mathrm{COST}(S)$$

Band we can restrict ourselves to periodic schedules. The proof is then along the lines of [5]. More precisely, if S is a periodic schedule of period T, we use Corollary 2. Given that $\sum_{j=1}^{n_i} t_j^i = T$, the sum of the squares is minimized when all terms are equal, to T/n_i. Defining $\tau_i = T/n_i$, we obtain the lower bound given in the first statement of the Lemma.

To solve the system, as in [5] 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 objective function, and $g(\tau) = C - \sum_{i=1}^m \frac{\ell_i}{\tau_i}$ 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 $\sum_{i=1}^m \ell_i \sqrt{\frac{p_i}{2c_i}} < C$. 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 colinear, 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 $\lambda^* \geqslant 0$:

$$(\Sigma) \left\{ \begin{array}{l} \mathrm{grad} f(\tau^*) = \lambda \ \mathrm{grad} g(\tau^*) \\ g(\tau^*) = 0 \end{array} \right.$$

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

$$\begin{array}{lll} (\Sigma) & \iff & (\tau^*, \lambda) \text{ is a critical point of } F \\ & \iff & \left\{ \begin{array}{l} (\forall i) \dfrac{\partial F}{\partial \tau_i}(\tau^*, \lambda) = 0 \\ \dfrac{\partial F}{\partial \lambda}(\tau^*, \lambda) = 0 \end{array} \right. \end{array}$$

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

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

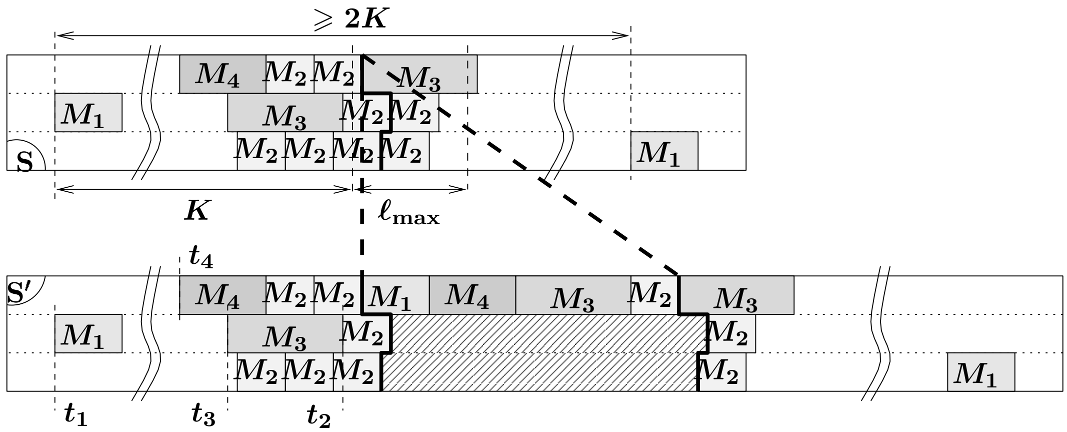

Figure 2: An example of the construction of S' from S in Lemma 4

$\lambda^* \geqslant 0$ satisfying $\sum_{i=1}^m \ell_i \sqrt{\frac{p_i}{2c_i + \lambda^* \ell_i}} = C$.

We have shown that f admits a unique minimum τ^* on domain D. Then a periodic schedule whose cost is equal to LB_0, necessarily broadcasts M_i exactly every τ_i^*. □

The corresponding lower bound was tight in the setting of [5]. However, in our model where the message lengths vary, this first lower bound is unsatisfying, since $\mathrm{OPT}/\mathrm{LB}_0$ can be arbitrarily large, as shown on the following example:

Example ▷ Consider the problem of scheduling $C+1$ messages $M_1, \ldots, M_{C+1}$ on C channels. M_i has length $\ell_i = L^{i-1}$, cost $c_i = 0$, and request probability $p_i = 1/L^{i-1}$ if $i > 1$, and $p_1 = 1 - \sum_{i=2}^m p_i$. In that case, it can be shown by induction on C that as L goes to infinity: $\mathrm{OPT} = \Theta(L^{1/2^C})$ but $\mathrm{LB}_0 = \Theta(1)$. ◁

Upon close examination, one realizes that the cost of a schedule in this example is unavoidably high even when $C = 1$ because, whenever a really long message is broadcast, all incoming requests have to wait. This observation led us to the following improved lower bound for the one-channel case. We will see in Theorem 10 that this improved lower bound is tight up to a constant factor in the one-channel case.

Remark ▶ Note that the cost of an optimal schedule on $C+1$ channels of these $(C+1)$ messages is $\Theta(1) \ll \Theta(L^{1/2^C})$. This emphasizes a big difference with the uniform length case where the cost is only reduced by a factor of $C/(C+1)$ from C to $C+1$ channels. ◀

Lemma 6 ($\mathbf{LB_\alpha}$) *Let $0 < \alpha < 1$. The following minimization problem is a lower bound to the cost of a schedule of $M_1, \ldots, M_m$ on one-channel:*

$$\mathrm{LB}_\alpha \left\{ \begin{array}{c} \min_{\tau > 0} \frac{1}{2} \sum_{1 \leqslant i \leqslant m} (1-\alpha) p_i \tau_i + \alpha \frac{\ell_i^2}{\tau_i} + \frac{2c_i}{\tau_i} \\ \text{subject to} \sum_i \frac{\ell_i}{\tau_i} \leqslant 1 \end{array} \right.$$

This minimization problem admits a unique solution τ^* verifying: $\tau_i^* = \sqrt{\frac{2c_i + \alpha \ell_i^2 + \lambda^* \ell_i}{(1-\alpha) p_i}}$ for a certain $\lambda^* \geqslant 0$. Furthermore, this lower bound is realized by a periodic schedule S if and only if M_i is scheduled in S periodically at intervals of lengths exactly τ_i^*.

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

Proof. As before we restrict ourself to periodic schedules (on one channel). 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 $\tau_i = \frac{T}{n_i}$. According to the proof above we have:

$$\mathrm{ART}(S) \geqslant \frac{1}{2} \sum_{i=1}^m p_i \tau_i \text{ and } \mathrm{BC}(S) = \sum_{i=1}^m \frac{c_i}{\tau_i}$$

We will show here another lower bound on $\mathrm{ART}(S)$. Consider a random request. With probability $\frac{n_i \ell_i}{T} = \ell_i/\tau_i$, it arrives during the broadcast of M_i, and then has to wait at least until M_i is finished broadcasting, i.e. $\ell_i/2$ on average. We conclude that:

$$\mathrm{ART}(S) \geqslant \sum_{i=1}^m \frac{\ell_i}{2} \frac{\ell_i}{\tau_i} = \frac{1}{2} \sum_{i=1}^m \frac{\ell_i^2}{\tau_i}$$

A linear $(1-\alpha, \alpha)$ combination of the two lower bounds on the average response time $\mathrm{ART}(S)$ gives us LB_α. The remaining statements in the lemma are immediatly deduced from the proof of Lemma 6 by substituting $c_i + \alpha\ell_i^2/2$ for c_i, and $(1-\alpha)p_i$ for p_i. □

4 On finding an optimal solution

4.1 *NP*-hardness of the Data-Broadcast problem

Theorem 7 is the first major result of this paper. Here, we restrict ourselves to the case when there are no broadcast costs. We show that even in this restrictive setting, and even when there is only one channel, the Data Broadcast problem is already *NP*-hard.

Theorem 7 (*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 on one channel is NP-hard.*

Proof. Based on a reduction from N-Partition (See Appendix page 11). □

4.2 The single channel problem with no broadcast costs

In the rest of this section, we assume that all the messages have zero broadcast cost and focus on the single channel case. Since no broadcasts cost are assumed, it never helps for a schedule to remain idle. Consequently, we can restrict our search for optimal schedules to schedules "without holes" (never idle). (As we will see in some examples at the end of section, this restriction is not valid in a more general setting).

In this special setting, the situation is much better understood, as demonstrated by the following results. The first lemma states that in this situation, there is an optimal schedule which is periodic.

Lemma 8 (Periodicity of an opt. schedule) *If the messages have no broadcast costs, then there exists an optimal schedule S^* of $M_1,\dots,M_m$ on one channel, which is periodic.*

Proof. We use Lemma 3 again. We will prove that among periodic schedules, there exists a schedule S^* of minimum cost. Then

$$\begin{aligned}\mathrm{COST}(S^*) &= \inf_{S \text{ periodic}} \mathrm{COST}(S)\\ &= \inf_S \mathrm{COST}(S) = \mathrm{OPT}\end{aligned}$$

and so S^* is optimal among all schedules.

Moreover, thanks to Lemma 4, we can restrict our search to periodic schedules without holes where any interval and where M_i is not broadcast has length at most K_i. Let us call such schedules *periodic bounded interval schedules without holes*.

The existence of an optimal schedule S^* among periodic schedules, relies on a graph construction similar to [3, 4]. We consider the following weighted labelled directed infinite graph G with costs. A vertex of G is a m-tuple $\langle a_1,\dots,a_m\rangle$ with $0 \leqslant a_i \leqslant K_i$. There is an edge from $\langle a_1,\dots,a_i,\dots,a_m\rangle$ to every vertex of the form $\langle a_1+\ell_i,\dots,\ell_i,\dots,a_m+\ell_i\rangle$. The *label* of this edge is M_i, its *length* is ℓ_i and its *cost* is $c(e) = \frac{1}{2}(p_i l_i^2 + \sum_{j=1,\ j\neq i}^{m} p_j\{(a_j+\ell_i)^2 - (a_j)^2\}) = \ell_i^2/2 + \sum_{j=1,\ j\neq i}^{m} p_j a_j \ell_i$. Semantically, being at node $u = \langle a_1,\dots,a_m\rangle$ at time t means that for all i, the last broadcast of M_i occured at time $t-a_i$; following an edge e with label M_i means that M_i is currently broadcast; then the waiting time for $M_{j\neq i}$ increases by ℓ_i and the waiting time for M_i is reduced to ℓ_i; the extra-cost induced by the broadcast of M_i is, according to Lemma 1, $c(e)$; and the extra-length of the schedule is ℓ_i.

More precisely, there is a natural surjection from cycles of G passing through a vertex $\langle 0, a_2,\dots,a_m\rangle$ to the finite-cost periodic schedules where any interval where M_i is not broadcast has length less than K_i. To each cycle γ going through $\langle 0,a_2,\dots,a_m\rangle$, we can associate the periodic schedule $S(\gamma)$ which is the sequence of labels of the edges of γ. This is clearly surjective. Furthermore, our definition of edge cost insures that:

$$\mathrm{COST}(S(\gamma)) = \frac{\sum_{e\in\gamma} c(e)}{\mathrm{length}(\gamma)} \underset{\text{def}}{=} \mathrm{COST}_G(\gamma)$$

Finally, consider a vertex $\langle a_1,\dots,a_m\rangle$ on cycle γ. Every coordinate a_i is a non-negative linear combination of ℓ_j's which sums to at most K_i, thus there are only a finite number of such vertices, thus only a finite number of such cycles. Thus, thanks to the surjection, there are only a finite number of periodic bounded interval schedules without holes. Then there is an optimal cost schedule among them, which is optimal among all the schedules. □

An important point is that, unlike [5], this lemma is not longer true if we consider schedule on more than one channel.

Example ▷ 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. Then one can show that $\mathrm{LB}_0(M_1, M_2) = \frac{\sqrt{2}}{2(1+\sqrt{2})}$, and this lower bound is reached by the schedule which never stops broadcasting M_1 on the first channel and M_2 on the second. If there was a periodic optimal schedule, it should, according to Lemma 5, broadcast M_1 every $\tau_1^* = 1$ time and M_2 every $\tau_2^* = \sqrt{2}$ time: impossible since $1/\sqrt{2}$ is irrational. ◁

In fact, idle periods have to exist in the presence of broadcast costs (for instance, a single message with a large cost should not be scheduled too often), or if more than one channel is available: it can be shown that an optimal schedule over two channels for two messages M_1 and M_2 with lengths $\ell_1 = 2$ and $\ell_2 = 3$, request probabilities $p_1 = 1 - \varepsilon$ and $p_2 = \varepsilon$, and zero broadcast cost, has to be idle after each broadcast of M_2 to allow synchronization. We do not currently know how to take the idle periods into account in the graph construction.

As a consequence of the proof of this last Lemma, we have the theorem below.

Theorem 9 *There exists a finite-time algorithm for constructing the optimal schedule for the data-broadcast messages of zero broadcast costs on one channel.*

Proof. According to the construction of Lemma 8, the algorithm just consists in finding a cycle γ^* of minimum cost in the subgraph of G restricted to the reachable vertices, and reading off the labels of the edges of γ^*. Note that in case of integer (or rational) message lengths, the algorithm of [19] for searching the minimum mean cycle can be used, hence an exponential time complexity of $\Theta\left(m(\ell_{\max}\mathcal{L}(4+3/p_{\min}))^{2m+1}\right)$. □

5 A single channel constant factor approximation algorithm

We now turn back to the more general setting where messages have non-zero broadcast costs. The theorem below is the second major result in this paper. It gives a randomized $O(1)$-approximation algorithm for the Data Broadcast problem on one channel. We denote by $E[x]$ the expected value of the random variable x.

Theorem 10 *There exists a randomized polynomial-time algorithm which, given α and ε, ouputs a schedule S on one channel satisfying:*

$$E[\mathrm{COST}(S)] \leqslant \max\left(\tfrac{1}{\alpha}, \tfrac{2}{1-\alpha}\right) \cdot \mathrm{LB}_\alpha(M_1, \ldots, M_m) + \varepsilon$$

For $\alpha = \frac{1}{3}$, this translates into:

$$E[\mathrm{COST}(S)] \leqslant 3\ \mathrm{LB}_{\alpha=\frac{1}{3}} + \varepsilon$$

Thus the lower bound designed in Lemma 6 is tight:

Corollary 11 $\mathrm{OPT} \leqslant 3\,\mathrm{LB}_{\alpha=\frac{1}{3}}$

Remark ▶ Thanks to the law of the large numbers, one can prove that the expected value of $\mathrm{COST}(S)$ is obtained with probability 1. ◀

Proof. Design of the algorithm. The algorithm is based on the lower bound proved in Lemma 6. This lower bound is quite informative since it can be used to calculate the desirable frequency $n_i/T = 1/\tau_i^*$ of each message M_i. Our algorithm will use randomness to produce a schedule according to these frequencies.

We now assume that we have managed to compute a good approximation $\tilde{\tau}_i$ of τ_i^*, so that $\sum_{i=1}^m \ell_i/\tilde{\tau}_i \leqslant 1$. As in the proof of Lemmas 5 and 6, two cases occur:

- if $\sum_{i=1}^m \ell_i/\tilde{\tau}_i = 1$: this is the easy case. Consider a schedule S which broadcasts each M_i with frequency $1/\tilde{\tau}_i$. The proportion of time spent broadcasting M_i is $\ell_i/\tilde{\tau}_i$. Then, in that case, S never stops broadcasting messages: there is no point in ever remaining idle.
- if $\sum_{i=1}^m \ell_i/\tilde{\tau}_i < 1$: this means that a schedule which broadcasts each M_i with frequency $1/\tilde{\tau}_i$ relative to the broadcasts of the messages should sometimes stop broadcasting messages and remain idle during a time proportional to $1 - \sum_{i=1}^m \ell_i/\tilde{\tau}_i$. Since we want the description of the schedule to remain discrete and not continuous, we introduce an artificial "ghost" message M_0 whose length ℓ_0 is a parameter of our algorithm. Its frequency $1/\tilde{\tau}_0$ will be adjusted as a function of ℓ_0 so that its broadcasts correspond to the idle times proposed by the lower bound. Thus $\tilde{\tau}_0$ is chosen so that:

  $$\ell_0/\tilde{\tau}_0 = 1 - \sum_{i=1}^m \ell_i/\tilde{\tau}_i > 0$$

  Now, consider a schedule which broadcasts each M_i with frequency $1/\tilde{\tau}_i$. It never stops broadcasting messages. The time periods spent in broadcasting M_0 correspond to required idle time periods.

In the first case the ghost message M_0 was not needed. We introduce it for the sake of homogeneity, with a frequency $1/\tilde{\tau}_0 = 0$.

Description of the algorithm. The algorithm works in 3 steps.

1. We compute $\tilde{\lambda}$, an estimation of the desired λ^* within ε accuracy, and the corresponding values $\tilde{\tau}_i$ such that $\tilde{\tau}_i \approx \sqrt{\frac{2c_i+\alpha\ell_i^2+\tilde{\lambda}\ell_i}{(1-\alpha)p_i}}$. This is easily done by dichotomy, as suggested in [5]. We then adjust $\tilde{\tau}_0$ so that $\ell_0/\tilde{\tau}_0 = 1 - \sum_{i=1}^m \ell_i/\tilde{\tau}_i$.
2. We compute the schedule frequency s_i of each message M_i, for $i = \mathbf{0}..m$: $$s_i = \frac{1/\tilde{\tau}_i}{\sum_{j=\mathbf{0}}^m 1/\tilde{\tau}_j}$$
3. We construct the schedule on-the-fly in a randomized fashion. Whenever the sender is ready to broadcast a message, with probability s_0 it remains idle for an interval of length ℓ_0, and with probability $s_{i\geqslant 1}$ it broadcasts M_i.

Analysis of the performance ratio. A short calculation shows that if:

$$\digamma(\lambda) = \frac{1}{2}\sum_{1\leqslant i\leqslant m}(1-\alpha)p_i\tau_i(\lambda) + \alpha\frac{\ell_i^2}{\tau_i(\lambda)} + \frac{2c_i}{\tau_i(\lambda)}$$

with $\tau_i(\lambda) = \sqrt{(2c_i+\alpha\ell_i^2+\lambda\ell_i)/(1-\alpha)p_i}$, then we have $0 \leqslant F'(\lambda) \leqslant 1/2$ whenever $\lambda \geqslant \lambda^* \geqslant 0$. Thus if our estimation $\tilde{\lambda}$ is close to λ^*, the cost function $\digamma(\tilde{\lambda})$ will also be close to the optimum:

$$\digamma(\tilde{\lambda}) \leqslant \digamma(\lambda^*) + \frac{1}{2}(\tilde{\lambda}-\lambda^*) \qquad \text{for } \tilde{\lambda} \geqslant \lambda^*$$

In particular if $\tilde{\lambda}$ approximates λ^* from above within ε, then $\digamma(\tilde{\lambda}) \leqslant \digamma(\lambda^*) + \varepsilon/2$.

We will now analyze the cost of the schedule obtained in step 3 of the algorithm.

Lemma 12 (Analysis of randomized sched.) *Consider the schedule S generated by the randomized process described at step 3 of the algorithm. Then:*

$$E[\mathrm{COST}(S)] = \frac{1}{\sum_{j=0}^m s_j\ell_j}\cdot\frac{1}{2}\sum_{i=0}^m s_i\ell_i^2 + \sum_{i=1}^m p_i\,\frac{\sum_{j=0,\ j\neq i}^m s_j\ell_j}{s_i} + \sum_{i=1}^m c_i\frac{s_i}{\sum_{j=0}^m s_j\ell_j}$$

Before proving the lemma, let us first see how it implies the claimed performance ratio.

Since s_i is proportional to $1/\tilde{\tau}_i$ and $\sum_{i=0}^m \ell_i/\tilde{\tau}_i = 1$, we get that:

$$\begin{aligned} E[\mathrm{COST}(S)] &= \tfrac{1}{2}\sum_{i=0}^m \tfrac{\ell_i^2}{\tilde{\tau}_i} + \sum_{i=0}^m p_i\tilde{\tau}_i\underbrace{(1-\ell_i/\tilde{\tau}_i)}_{0\leqslant\ \ \leqslant 1} + \sum_{i=0}^m \tfrac{c_i}{\tilde{\tau}_i} \\ &\leqslant \tfrac{\ell_0^2}{2} + \max\left(\tfrac{1}{\alpha}, \tfrac{2}{1-\alpha}, 1\right)\digamma(\tilde{\tau}) \end{aligned}$$

As the minimum of $\max(\frac{1}{\alpha}, \frac{2}{1-\alpha}, 1)$ is 3, obtained for $\alpha = 1/3$ then we get, for $\alpha = 1/3$:

$$E[\mathrm{COST}(S)] \leqslant 3\,\digamma(\tilde{\tau}) + \frac{\ell_0^2}{2} \leqslant 3\,\mathrm{OPT} + \frac{\varepsilon+\ell_0^2}{2}$$

□

Remark ► Taking $\varepsilon \leqslant \eta\digamma(\tilde{\tau})$ and $\ell_0 \leqslant \sqrt{\eta\digamma(\tilde{\tau})}$ yields the absolute performance ratio of $3+\eta$. ◄

Let us now turn to the proof of the lemma.

Proof. Consider M_i, $1 \leqslant i \leqslant m$. Let us define the random variable $R_i(t)$ as the response time to a request for M_i arriving at time t. The average response time is the weighted sum of $E(R_i(t))$, so we just need to compute the expectation of $R_i(t)$.

We decompose $R_i(t)$ into two random variables $R_i(t) = X(t) + Y_i$: $X(t)$ is the time until the end of the message which is broadcast at time t, and Y_i is the time from then until M_i is broadcast; note that $X(t)$ is independent of i and that Y_i is independent of t.

We first compute $E[X]$. If the request is raised during the broadcast of M_j, it will wait on average $\ell_j/2$. Thus:

$$E[X] = \sum_{j=0}^m \frac{\ell_j}{2}\cdot\Pr\left\{\begin{array}{l}\text{request raised during}\\ \text{broadcast of } M_j\end{array}\right\}$$

But, $\Pr\{\text{request raised during } M_j\}$
$= \ell_j \cdot E\big[\limsup_{T\to\infty}\frac{N_j(T)}{T}\big]$
$= \ell_j \cdot \int_S \limsup_{T\to\infty}\frac{N_j(T)}{T}\, d\Pr\{S\}$

where $N_j(T)$ is the number of broadcasts of M_j during the time period $[0,T]$. Since $T \geqslant N_j(T)\ell_j$, Lebesgues's bounded convergence theorem implies that $\Pr\{\text{request raised during } M_j\}$
$= \ell_j \cdot \limsup_{T\to\infty} E\big[\frac{N_j(T)}{T}\big]$

But $E\big[\frac{N_j(T)}{T}\big] \sim_{T\to\infty} \frac{s_j}{\sum_{k=0}^m s_k\ell_k}$, thus $\Pr\{\text{request raised during } M_j\} = \frac{s_j\ell_j}{\sum_{k=0}^m s_k\ell_k}$. Finally:

$$E[X] = \frac{1}{\sum_{k=0}^m s_k\ell_k}\sum_{j=0}^m \frac{s_j\ell_j^2}{2}$$

Let us now compute $E[Y_i]$. We define Z_q as the length of the qth message broadcast by the algorithm, conditioned on being different from M_i. We have then:

$$E[Y_i] = E\left[\sum_{q\geqslant 0}(Z_1+\cdots+Z_q)(1-s_i)^q s_i\right]$$

But $E[Z_1] = \ldots = E(Z_q) = \sum_{j=0,\ j\neq i}^m \frac{s_j}{1-s_i}\ell_j$, thus an algebraic manipulation yields:

$$E[Y_i] = \frac{\sum_{j=0,\ j\neq i}^{m} s_j \ell_j}{s_i}$$

hence $E(R_i)$ and $E(\text{ART}(S))$.

Finally $\text{BC}(S, [0, T]) = \sum_{i=1}^{m} c_i N_i(T)/T$. Since the expected frequence of M_i is $E\big[\limsup_{T\to\infty} \frac{N_i(T)}{T}\big] = \frac{s_i}{\sum_{j=0}^{m} s_j \ell_j}$, we get:

$$E[\text{BC}(S)] = \sum_{i=1}^{m} \frac{c_i s_i}{\sum_{j=0}^{m} s_j \ell_j}$$

which finishes the proof of the lemma. □

References

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

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

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

[4] Anily, S., Glass, C. A., and Hassin, R. The scheduling of maintenance service. In *Discrete Applied Math.* to appear.

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

[6] Chan, M. Y., and Chin, F. General schedulers for the pinwheel problem based on double-integer reduction. In *IEEE Transaction on Computers* (Jun. 1992), vol. 41(6), pp. 755–768.

[7] Chan, M. Y., and Chin, F. Schedulers for larger classes of pinwheel instances. *Algorithmica 9* (Jun. 1993), 425–462.

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

[9] Hadley, G., and Whitin, T. M. *Analysis of inventory systems.* Prentice-Hall, 1963.

[10] Hofri, M., and Rosberg, Z. Packet delay under the golden ratio weighted TDM policy in a multiple-access channel. In *IEEE Trans. on Information Theory* (May 1987), vol. 11(33), pp. 341–349.

[11] Holte, R., Mok, A., Rosier, L., Tulchinsky, I., and Varvel, D. The pinwheel: A real-time scheduling problem. In *Proc. 22nd Hawaii Int. Conf. Syst. Sci.* (Jan. 1989), pp. 693–702.

[12] Itai, A., and Rosberg, S. A golden ratio control policy for a multiple-access channel. *IEEE Transactions on Autom. Contr. AC-29* (Aug. 1984).

[13] Khanna, S., and Liberatore, V. On broadcast disk paging. In *Proceedings of 30th STOC* (1998), vol. 30, pp. 634–643.

[14] Khanna, S., and Zhou, S. On indexed data broadcast. In *Proceedings of 30th STOC* (1998), vol. 30, pp. 463–472.

[15] Knuth, D. *The art of computer programming, Sorting and Searching*, vol. 3. 1973.

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

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

[18] Vaidya, N., and Hameed, S. 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).

[19] Young, N. E., Tarjan, R. E., and Orlin, J. B. Faster parametric shortest path and minimum balances algorithms. In *Networks* (Mar. 1991), vol. 21(2).

Appendix A: Omitted proofs

Lemma 3 (Density of periodic schedules) *Let S be a schedule of $M_1, \ldots, M_m$ on C channels. For any $\varepsilon > 0$, there exists a* periodic *schedule S' whose cost satisfies:*

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

Proof sketch. If the cost of S is infinite, it is sufficient to pick for S' any periodic schedule in which all messages are broadcast at least once. The interesting part of the proof occurs when the cost of S is finite.

Construction of S'. Let T be such that all messages have been broadcast at least once before T and such that for every $t \geqslant T$, we have:

$$\begin{cases} \mathrm{ART}(S, [0,t]) & \leqslant & \mathrm{ART}(S) + \varepsilon/4 \\ \mathrm{BC}(S, [0,t]) & \leqslant & \mathrm{BC}(S) + \varepsilon/2 \end{cases}$$

Consider $t \geqslant T$, whose exact value will be given later; let S_t denote the schedule obtained from S by removing all messages in the process of being broadcast at time t, and let S' denote the unique periodic schedule of period t which is identical to S_t over the time interval $[0, t]$.

Analysis of S'. Clearly, the average response time of S' over $[0, t]$ is the same as the average response time of S over $[0, t]$, except for boundary effects. In fact, this can be quantified using Lemma 1, and one can then show a bound of the following form:

$$\mathrm{ART}(M_i, S') \leqslant \mathrm{ART}(M_i, S, t) + \frac{c}{\sqrt{t}}$$

where c is a constant which only depends on $\mathrm{ART}(S)$, t_1^i, p_i, and ε. Choosing $t = \max\{T, 16/c^2\varepsilon^2\}$ then yields the result of the lemma. □

Lemma 4 (Maximum interval) *Let S be a periodic schedule of $M_1, \ldots, M_m$ on C channels. There exists a periodic schedule S' such that for any i, any interval where M_i is not broadcast has length at most K_i, where:*

$$K_i = \frac{3\mathcal{L}}{p_i} + 4\mathcal{L} + \frac{2\,\mathcal{C}}{\mathcal{L}},\ \textit{with}\ \mathcal{L} = \sum_{j=1}^{m} \ell_j,\ \mathcal{C} = \sum_{j=1}^{m} c_j$$

and such that: $\mathrm{COST}(S') \leqslant \mathrm{COST}(S)$

Proof sketch. Detailed construction of S'. The construction is illustrated on Figure 2. Let $\ell_{\max}$ denote the maximum message length, i.e. $\ell_{\max} = \max_i \ell_i$. From S, we construct a periodic schedule S' of period $T + \mathcal{L}$ as follows. For every channel k, $1 \leqslant k \leqslant C$, let T_k be equal to the time when channel k finishes broadcasting the message being broadcast at time K if there is one, and equal to K otherwise. Note that $K \leqslant T_k \leqslant K + \ell_{\max}$. Suppose, w.l.o.g., that $T_1 = \min_k T_k$.

S' is obtained from S by doing the following: first, for every $1 \leqslant k \leqslant C$, introducing an hole of size $\mathcal{L}$ at time T_k on channel k; second, inserting on channel 1 all the messages one after the other in order of increasing t_i, where, for all $1 \leqslant i \leqslant m$, we define t_i as the date of the beginning of the last broadcast of M_i on any channel before T_1. This guarantees that the intervals of time between two broadcasts of M_i, $i \geqslant 2$, do not change by much.

Analysis of S'. Let us show that $\Delta(T\,\mathrm{COST}(S)) = (T + \mathcal{L})\,\mathrm{COST}(S') - T\,\mathrm{COST}(S) \leqslant 0$. Since $T + \mathcal{L} > T$, the lemma will follow.

Split the messages $M_2, \ldots, M_m$ into two sets I and J. In the first set, we put the messages for which the insertion of all the messages in S' may be bad: that is to say $M_i \in I$ if $i \geqslant 2$ and the time elapsed between the last broadcast of M_i before T_1 and the next one, was less than K in S. We define J as $J = \{M_2, \ldots, M_m\} \setminus I$. One can show then the following upper bounds:

$$\Delta(T\,\mathrm{COST}(S))_{\text{restricted to } M_i \in I} \leqslant \textstyle\sum_{i \in I} p_i \mathcal{L}(K + \mathcal{L}) + c_i$$

$$\Delta(T\,\mathrm{COST}(S))_{\text{restricted to } M_j \in J} \leqslant \textstyle\sum_{j \in J} p_j \mathcal{L}(K + \ell_{\max}) + c_j$$

$$\Delta(T\,\mathrm{COST}(S))_{\text{restricted to } M_1} \leqslant -p_1(K - \mathcal{L})^2 + p_1\mathcal{L}(K + \ell_{\max}) + c_1$$

Thus, $\Delta(T\,\mathrm{COST}(S)) \leqslant -p_1(K - \mathcal{L})^2 + \sum_{i=1}^{m} p_i \mathcal{L}(K + \mathcal{L}) + \mathcal{C}$

¿From there, it is easy to see that if $K \geqslant \frac{1}{2}(\mathcal{L}(4 + \frac{3}{p_1}) + \frac{2\,\mathcal{C}}{\mathcal{L}})$, then $\Delta(T\,\mathrm{COST}(S)) \leqslant 0$.

Applying this scheme recursively to each pair of consecutive occurrences of M_1 completes the proof. □

Theorem 7 (*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 on one channel is NP-hard.*

Proof. We will reduce N-Partition to the Data-Broadcast problem without costs as follows. Recall that N-Partition takes as input m integers $x_1, \ldots, x_m$ and must decide whether there exists a partition of $\{1, \ldots, m\}$ in N disjoint sets $I_1, \ldots, I_N$ such that $\sum_{i \in I_1} x_i = \cdots = \sum_{i \in I_N} x_i$. Let $\mathcal{S} = x_1 + \cdots + x_m$.

We consider the following one-channel Data-Broadcast instance: $m+1$ messages $M_0, M_1, \ldots, M_m$ such that:

$$\begin{cases} M_0: & \ell_0 = \mathcal{S}/N^2 \text{ and } p_0 = 1/2 \quad (c_0 = 0) \\ M_{i \geqslant 1}: & \ell_i = x_i \text{ and } p_i = \dfrac{x_i}{2\mathcal{S}} \quad (c_i = 0) \end{cases}$$

In that case, one can prove that $\mathrm{LB}_0(M_0, \ldots, M_m) = \mathcal{S}/4 \cdot (1 + \frac{1}{N})^2$. We consider the decision problem "Does there exist a schedule with cost less than or equal to $\mathcal{S}/4 \cdot (1 + \frac{1}{N})^2$?".

Note that this reduction is clearly polynomial. According to Lemmas 5 and 8, the answer to the question is "yes" if and only if there exists a periodic schedule which broadcasts each M_i every τ_i^* exactly. One can show that here: $\tau_0^* = \mathcal{S}/N + \ell_0$ and $\tau_{i \geqslant 1}^* = \mathcal{S} + N\ell_0$. Such a schedule must then be of the following form: between two consecutive broadcasts of M_0 a sequence of messages of total length at most $\mathcal{S}/N$ can be broadcast. Then the answer is "yes" if and only if the set of messages $M_1, \ldots, M_m$ can be split into N sets $\mathcal{M}_1, \ldots, \mathcal{M}_N$ such that $\sum_{M_i \in \mathcal{M}_1} \ell_i = \cdots = \sum_{M_i \in \mathcal{M}_N} \ell_i$. □

Remark ▶ This proof can easily be generalized to prove the *NP*-hardness of the problem "Does there exists a *periodic* schedule of m messages, with rational lengths and probabilities, and zero broadcast costs, on C channels whose cost is less or equal than x?". Taking the reduction above, it suffices to add $C - 1$ messages $M_{m+1}, \ldots, M_{m+C-1}$ of lengths $\ell_{i>m} = \mathcal{S} + 2\ell_0$ and probabilities adjusted so that $\tau_{i>m}^* = \mathcal{S} + 2\ell_0$. ◀

Appendix B: Brief analysis of the time complexity of the single-channel algorithm.

Let us estimate the computation time of $\tilde{\lambda}$. Two cases occur. If $\sum_{i=1}^m \ell_i/\tau_i^* < 1$, then $\lambda^* = 0$ and thus $\tilde{\lambda} = 0$. If $\sum_{i=1}^m \ell_i/\tau_i^* = 1$, then following [5], we get that:

$$\begin{aligned} \min_i((m^2(1-\alpha)p_i - 1)\ell_i - \tfrac{c_i}{\ell_i}) &\leqslant \lambda^* \\ &\leqslant \max_i((m^2(1-\alpha)p_i - 1)\ell_i - \tfrac{c_i}{\ell_i}) \end{aligned}$$

Thus a $O(\log\{p_i, \ell_i, c_i/\ell_i, 1/\varepsilon, m\})$-time dichotomy gives an approximation $\tilde{\lambda} \geqslant \lambda^*$ of λ within ε accuracy. Since $\tilde{\lambda} \geqslant \lambda^*$ and $\tau^*(\lambda)$ is an increasing function, we ensure that $\sum_{i=1}^m \ell_i/\tau_i^*(\tilde{\lambda}) \leqslant 1$. A $O(\log\{p_i, \ell_i, c_i/\ell_i, 1/\varepsilon, m\})$-time square root computation gives then the desired approximation $\tilde{\tau}_i \geqslant \tau^*(\tilde{\lambda})$ within ε accuracy. Thus the $\tilde{\tau}_i$ are obtained by a linear time algorithm in the size of the entries.

Appendix C: A special case resolved

An exhaustive case-by-case analysis solves the problem completely in the single-channel zero-cost case when there are only two messages.

Theorem 13 *Let consider the case of the broadcast of two messages M_1 and M_2 with respective lengths, request probabilities and broadcast costs: ℓ_1 and ℓ_2, p_1 and p_2, $c_1 = c_2 = 0$. If $\beta = \ell_1/\ell_2$, then an optimal schedule has the following form:*

Condition	*Opt. Schedule*
$p_1 \leqslant \frac{\beta}{2(\beta+1)}$	$(M_1 M_2^{n_2})^\omega$, *where* $n_2 \geqslant 2$
$\frac{\beta}{2(\beta+1)} \leqslant p_1 \leqslant \frac{2\beta+1}{2(\beta+1)}$	$(M_1 M_2)^\omega$
$\frac{2\beta+1}{2(\beta+1)} \leqslant p_1$	$(M_1^{n_1} M_2)^\omega$, *where* $n_1 \geqslant 2$

where n_i is equal to either N_i or $N_i + 1$, with $N_i = \left\lfloor \dfrac{\ell_j}{\ell_i} \left(\sqrt{\frac{p_i}{p_j}\left(1 + \frac{\ell_i}{\ell_j}\right)} - 1 \right) \right\rfloor$, for $j \neq i$.