Time-Division is Better than Frequency-Division for Periodic Internet Broadcasting

Kevin Foltz Jehoshua Bruck

California Institute of Technology
Mail Code: 136-93
Pasadena, CA 91125

Email: {kfoltz,bruck}@paradise.caltech.edu

Abstract

The broadcast disk provides an effective way to transmit information from a server to many clients. Information is broadcast cyclically and clients pick the information they need out of the broadcast. An example of such a system is a wireless web service where web servers broadcast to browsing clients. We consider two ways to send items over a broadcast channel and compare them using the metric of expected waiting time. The first is frequency-division, where each item is broadcast on its own subchannel of lower bandwidth. We find the optimal allocation of bandwidth to the subchannels using this method. Then we look at time-division, where items are sent sequentially on a single full-bandwidth channel. For items of equal length, we show that for any frequency-division broadcast schedule, we can find a better time-division schedule. Thus time-division is better than frequency-division.

1 Introduction

As mobile computing gains popularity, it becomes increasingly important to find efficient methods of communicating with mobile clients. In general, the mobile clients have considerably less outgoing bandwidth than incoming bandwidth, making communication highly asymmetric. Web browsing is a good example of this situation. A person browsing the web typically receives a lot of information and sends relatively little. We compare frequency-division scheduling and time-division scheduling as ways to send information to portable web browsers. The measure of efficiency that we will use is "expected waiting time," as defined in the next section.

The broadcast disk is a way to send information to many clients at the same time. Using this scheme, data is repeatedly broadcast through the air. When a client wants some data, it listens to this broadcast until it receives the desired data. The broadcast disk acts as a common cache for many clients, where data in this cache is made available according to the broadcast schedule. The goal is to broadcast the information in a way that minimizes the expected waiting time of the clients.

The first broadcasting scheme we consider uses frequency-division. Using this scheme, we divide the total available bandwidth into channels and repeatedly broadcast each item on its own channel. We then examine time-division, where we allocate all bandwidth to a single channel and send items sequentially over this channel.

We find the optimal allocation of bandwidth to the subchannels for the frequency-division method based on the items' lengths and demand probabilities. We then show that for equal length items, given any frequency-division schedule, we can generate a time-division schedule with a lower expected waiting time.

Vaidya and Hameed [6, 7, 11] determined the optimal broadcast frequencies of items within a time-division schedule as a function of their demand probabilities, p_i, and lengths, l_i. They showed that to minimize expected waiting time, the items should be uniformly spaced, with frequencies of broadcast, f_i, proportional to $\sqrt{\frac{p_i}{l_i}}$. This led to an algorithm that attempted to achieve these relative frequencies. This algorithm is good because it is computationally fairly simple and works for an arbitrary number of broadcast items with arbitrary lengths and demand probabilities. However, in most cases it is not possible to achieve the optimal frequencies with uniform spacing. This can lead to expected waiting times that are worse than the theoretical optimal time. We looked at splitting items into pieces as a way to improve performance. We found the optimal schedules for two, equal-length items, when splitting them in half is allowed [4].

Jiang and Vaidya [8] look at ways to minimize the variance of the response time, and trade-off between minimizing the mean and variance of the response time. Kenyon and Schabanel [9] examine the broadcast of multiple items with different lengths and transmission costs. They show that finding optimal schedules is NP-Hard, and the optimal schedules seem very different structurally than schedules for equal lengths. Aksoy and Franklin [1] discuss scheduling the broadcast of information based on client requests. Bestavros [3] describes a way to add fault tolerance to broadcast disks by sending parity information in addition to data. Bar-Noy, Bhatia, Naor, and Schieber [2] look at scheduling in general, and show that there is an optimal cyclic schedule for a broadcast disk, and finding it is NP-hard. Leong and Si [10] discuss how to choose which items to broadcast, using ideas of cache management. Franklin, Zdonik, Alonso, and Acharya [5, 12] also discuss aspects of broadcast disks.

In the next section, we define some ideas related to schedules and expected waiting time. In Section 3, we derive the optimal frequency-division schedules. Then, in Section 4, we state our theorem regarding frequency-division and time-division, and present a proof. We conclude with

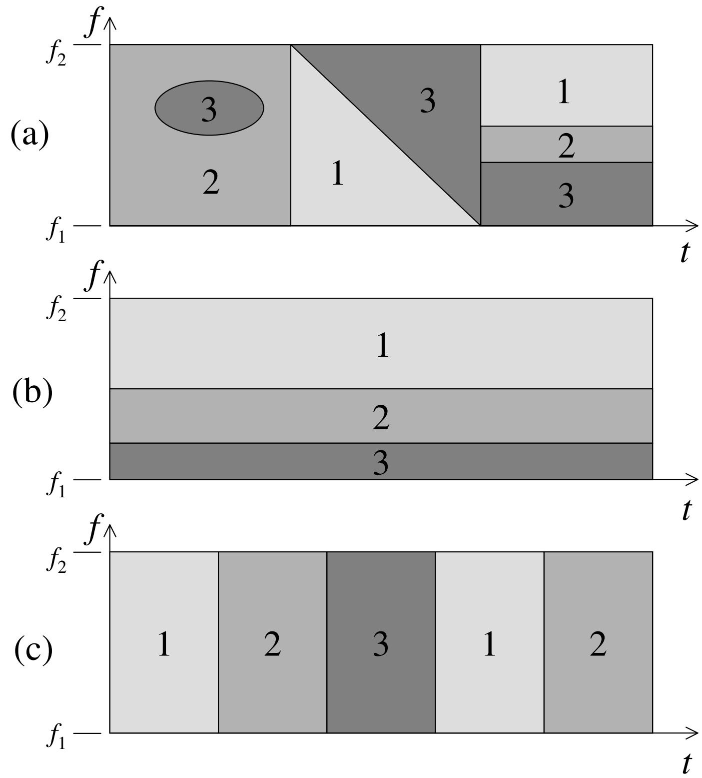

Figure 1: Examples of different types of schedules: (a) general schedule (b) frequency-division schedule, and (c) time-division schedule

Section 5, presenting some open problems and ideas for future work.

2 Types of Schedules

We define our notion of a schedule:

Definition 1 *A schedule for n items and frequency interval $[f_1, f_2]$ is a function $S : \mathbb{R} \times [f_1, f_2] \to \{1, 2, \ldots, n\}$, where $S(t, f)$ is the number of the item sent at time t over frequency f.*

Essentially, a schedule is an assignment of our two resources, time and bandwidth, to the n items we wish to broadcast. This definition is very generic, but we will only consider two special classes of schedules:

Definition 2 *A frequency-division schedule for n items and frequency interval $[f_1, f_2]$ is a schedule S for n items and frequency interval $[f_1, f_2]$ where, for all $t \in \mathbb{R}$ and $f \in [f_1, f_2]$, $S(t, f) = f_F(f)$, for some function $f_F : [f_1, f_2] \to \{1, 2, \ldots, n\}$.*

Definition 3 *A time-division schedule for n items and frequency interval $[f_1, f_2]$ is a schedule S for n items and frequency interval $[f_1, f_2]$ where, for all $t \in \mathbb{R}$ and $f \in [f_1, f_2]$, $S(t, f) = f_T(t)$, for some function $f_T : \mathbb{R} \to \{1, 2, \ldots, n\}$.*

Figure 1 represents these schedules graphically. We will generally consider periodic schedules S with $[f_1, f_2] = [0, B]$, where B is the bandwidth we have available for broadcast. We will refer to these as schedules of n items and bandwidth B. Using frequency-division, we essentially divide the bandwidth into n channels and repeatedly send one item per channel. With time-division, we simply send items one after another in some specified order using the full bandwidth available. To compare these two methods, we define "expected waiting time."

Our schedule tells us when we send each item, but does not tell us which part of the item to send. We assume that we send bits of each item sequentially within their bandwidth allocation. Since we must get an item from start to finish, the starting and ending points for the items are important to know. We define the ending point as follows:

Definition 4 *The k^{th} ending point for item i using schedule S, for $k \in \mathbb{Z}$, is the smallest time t for which $\int_0^t \int_{f_1}^{f_2} \delta\left(S\left(t,f\right),i\right) df\, dt = k \cdot l_i$*

Essentially, we assume that at time zero we are ready to send the first bit of each item, and the k^{th} ending point of item i is the place where we have just completed sending item i for the k^{th} time.

Definition 5 *The waiting time for item i of length l_i using schedule S (for n items and frequency interval $[f_1, f_2]$) and initial listening time t_1 is a function $WT_i(S, t_1) = t_2 - t_1 - \frac{l_i}{f_2 - f_1}$, where t_2 is the smallest ending point for item i such that $\int_{t_1}^{t_2} \int_{f_1}^{f_2} \delta(S(t,f), i) df\, dt \geq l_i$, where $\delta(a,b) = \begin{cases} 1 & \text{if } a = b \\ 0 & \text{otherwise} \end{cases}$.*

We subtract the term $\frac{l_i}{f_2 - f_1}$ because we will always listen at least this long for item i. Even in the best possible case, where we immediately start receiving item i from its beginning and continue receiving it over the full available bandwidth, we still listen a time $\frac{l_i}{f_2 - f_1}$. Since we are interested in studying waiting times due to scheduling, we subtract this constant "transmission time" when computing the waiting time.

Definition 6 *Expected waiting time (EWT) is a function of a schedule and a vector of demand probabilities $\vec{p} = (p_1, p_2, \ldots, p_n)$. $EWT\left(S, \vec{p}\right) = \sum_{i=1}^{n} p_i EWT_i(S)$, where $EWT_i(S) = \lim_{T \to \infty} \frac{1}{T} \int_0^T WT_i(S,t) dt$.*

Example 1

Figures 2 and 3 show sample WT computations for frequency-division and time-division schedules. Since each of the schedules is periodic with period $\frac{5}{B}$, we can compute $EWT_i(S)$ by simply computing the average value of $WT_i(S,t)$ for $t \in \left[0, \frac{5}{B}\right]$. For the frequency-division schedule, we get $EWT_1 = \frac{11}{4B}$, $EWT_2 = \frac{11}{4B}$, and $EWT_3 = \frac{13}{2B}$. For the time-division schedule, we get $EWT_1 = \frac{13}{10B}$, $EWT_2 = \frac{13}{10B}$, and $EWT_3 = \frac{5}{2B}$. For $p_1 = \frac{9}{20}$, $p_2 = \frac{7}{20}$, and $p_3 = \frac{1}{5}$, we get expected waiting times of $\frac{7}{2B}$ for the frequency-division schedule and $\frac{77}{50B}$ for the time-division schedule.

3 Optimal Frequency-Division Scheduling

We want to use frequency-division scheduling to send n items over a broadcast channel of bandwidth B. This is shown for $n = 3$ in Figure 4. When waiting for item i, a client listens to channel i of bandwidth $\alpha_i \cdot B$. We define a specific frequency-division schedule and prove its optimality:

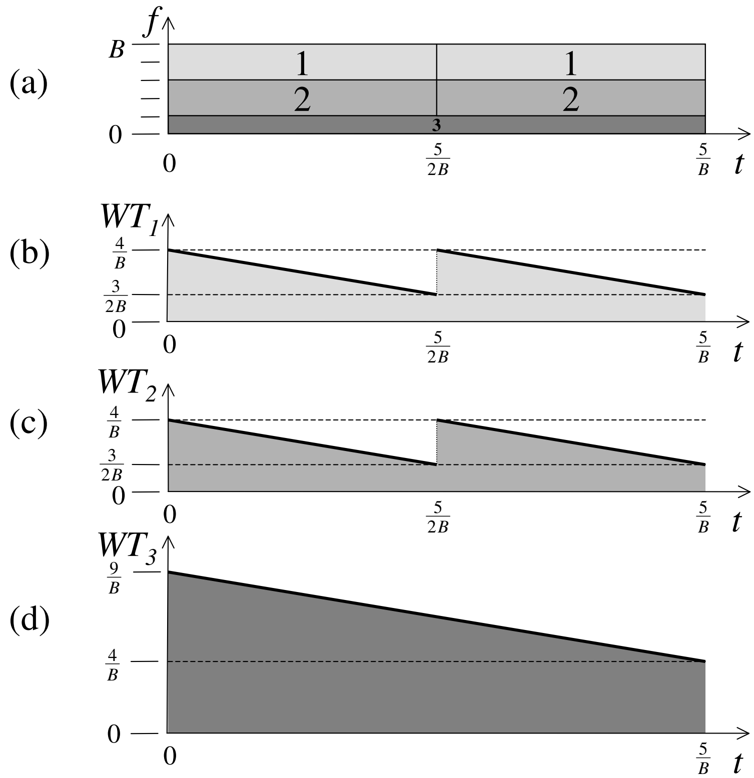

Figure 2: Sample WT calculation for a frequency-division schedule: (a) the schedule (b) WT_1 (c) WT_2, and (d) WT_3

Definition 7 *The square root rule frequency-division schedule (SRR-FD schedule) for n items of lengths l_i, $1 \leq i \leq n$, frequency interval $[f_1, f_2]$, and demand probability vector $\vec{p} = (p_1, p_2, \ldots, p_n)$ is a frequency-division schedule S for n items and frequency interval $[f_1, f_2]$, with $S(t, f) = k$, if $f_1 + (f_2 - f_1)\sum_{j=k+1}^{n} \alpha_j \leq k < f_1 + (f_2 - f_1)\sum_{j=k}^{n} \alpha_j$, where $\alpha_i = \frac{\sqrt{p_i l_i}}{\sum_{j=1}^{n} \sqrt{p_j l_j}}$, $1 \leq i \leq n$.*

Theorem 1 *Among all frequency-division schedules for n items of lengths l_i, $1 \leq i \leq n$, and frequency interval $[f_1, f_2]$, for demand probability vector $\vec{p} = (p_1, p_2, \ldots, p_n)$, the corresponding SRR-FD schedule has minimal expected waiting time.*

Proof of Theorem 1

First, we note that the waiting times depend only on the total bandwidth allocated to each item, and not how it is distributed in S. So, we are free to assume that any frequency-division schedule S consists of n different "subchannels", where item i is sent on subchannel i of bandwidth $\alpha_i B$ and $\sum_{i=1}^{n} \alpha_i = 1$. The waiting time, using the notation of Definition 5, is $WT_i = t_2 - t_1 - \frac{l_i}{f_2 - f_1}$. We see that $t_2 - t_1 = t_{partial} + t_{item}$, where t_{item} is the wait to receive the item from start to finish, starting at the next ending point, and $t_{partial}$ is the wait until the next ending point, which ranges from 0 to t_{item}. So $t_2 - t_1 = t_{partial} + \frac{l_i}{\alpha_i B}$, and $WT_i = t_{partial} + \frac{l_i}{\alpha_i B} - \frac{l_i}{B}$. We now compute $EWT_i = \lim_{T\to\infty} \int_0^T \left(t_{partial} + \frac{l_i}{\alpha_i B} - \frac{l_i}{B}\right) dt = \lim_{T\to\infty} \int_0^T t_{partial} dt + \frac{l_i}{\alpha_i B} - \frac{l_i}{B} = \frac{1}{2}\frac{l_i}{\alpha_i B} + \frac{l_i}{\alpha_i B} - \frac{l_i}{B} = \frac{3}{2}\frac{l_i}{\alpha_i B} - \frac{l_i}{B}$.

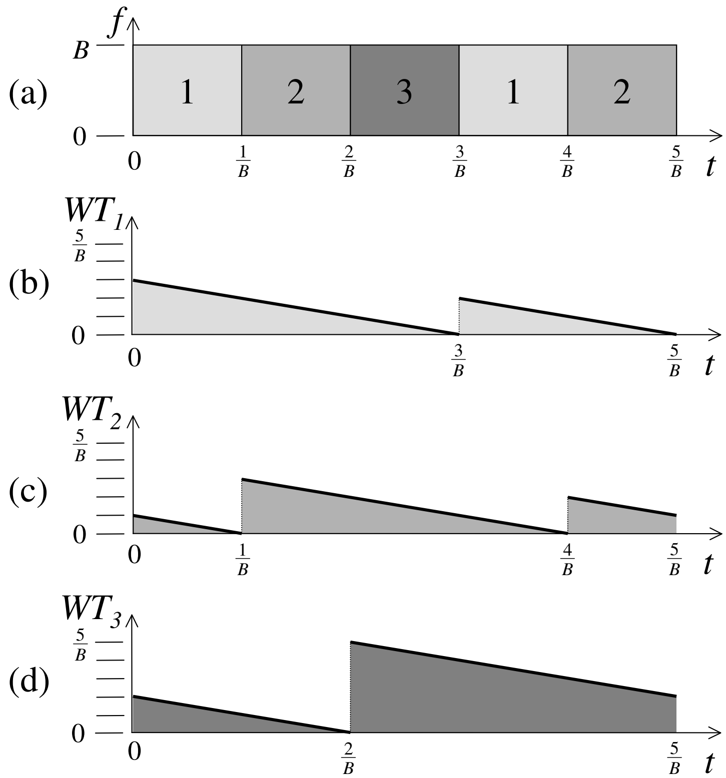

Figure 3: Sample WT calculation for a periodic time-division schedule: (a) one period of the schedule (b) WT_1 (c) WT_2, and (d) WT_3

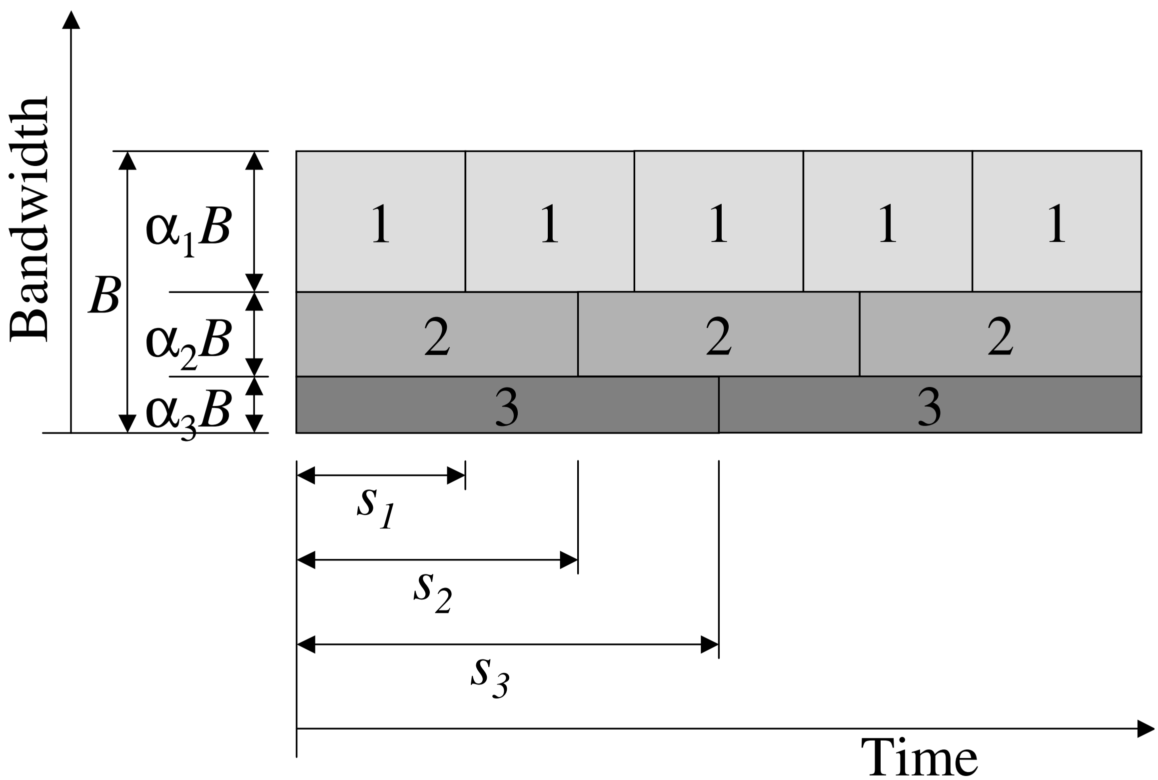

Figure 4: Frequency-division scheduling

This gives us an overall expected waiting time of:

$$
\begin{aligned}
EWT &= \sum_{i=1}^{n} \left(\frac{3}{2} \frac{l_i}{\alpha_i B} - \frac{l_i}{B} \right) p_i \\
&= \frac{3}{2B} \sum_{i=1}^{n} \frac{p_i l_i}{\alpha_i} - \frac{1}{B} \sum_{i=1}^{n} p_i l_i
\end{aligned}
$$

This is minimized when $\sum_{i=1}^{n} \frac{p_i l_i}{\alpha_i}$ is minimized, or when

$$
\alpha_i = \frac{\sqrt{p_i l_i}}{\sum_{j=1}^{n} \sqrt{p_j l_j}}.
$$

These are the values of the α_i's used by the SRR-FD schedule. It follows that the SRR-FD schedule has minimal expected waiting time among all frequency-division schedules.

Using these values of α_i, we get the minimal expected waiting time of:

$$
EWT_{min} = \frac{1}{B} \sum_{i=1}^{n} \sqrt{p_i l_i} \left(\frac{3}{2} \sum_{j=1}^{n} \sqrt{p_j l_j} - \sqrt{p_i l_i} \right)
$$

Example 2

As an example of frequency-division scheduling, we consider scheduling three items with lengths $l_1 = 9$, $l_2 = 32$, and $l_3 = 3$, and demand probabilities $p_1 = .46$, $p_2 = .05$, and $p_3 = .49$. We compute $p_1 l_1 = 4.14$, $p_2 l_2 = 1.6$, and $p_3 l_3 = 1.47$. Taking square roots and normalizing, we get $\alpha_1 = .45$, $\alpha_2 = .28$, and $\alpha_3 = .27$. Figure 5 shows the resulting schedule for $B = 1$.

4 Better Time-Division Scheduling

We now consider time-division schedules, as in Figure 6. Finding the optimal time-division schedules, as we did for frequency-division schedules, is difficult. Instead, we consider the case when all items have the same length, and show that for any frequency-division schedule we can find a time-division schedule that has lower expected waiting time.

Theorem 2 *For any frequency-division schedule, S_1, for n items of length l and frequency interval $[f_1, f_2]$, there is a time-division schedule, S_2, for n items of length l and frequency interval $[f_1, f_2]$ such that $EWT(S_2, \vec{p}) \leq EWT(S_1, \vec{p})$ for any vector $\vec{p} = (p_1, p_2, \ldots, p_n)$ of demand probabilities.*

Proof of Theorem 2

The idea of the proof is to consider an arbitrary frequency-division schedule and construct a better time-division schedule. We will assume WLOG that the length of the items is $l = 1$, the bandwidth is $B = 1$, and the α_i's are ordered as follows: $\alpha_1 \geq \alpha_2 \geq \alpha_3 \geq \ldots \geq \alpha_n$. We let $s_i = \frac{1}{\alpha_i}$, the spacing between instances of item i in the frequency-division schedule. For the frequency-division schedule, the expected waiting time is $\sum_{i=1}^{n} p_i \left(\frac{3}{2\alpha_i} - 1 \right)$. We will construct a time-division schedule that has a lower expected waiting time. We choose our construction based on the value of α_1.

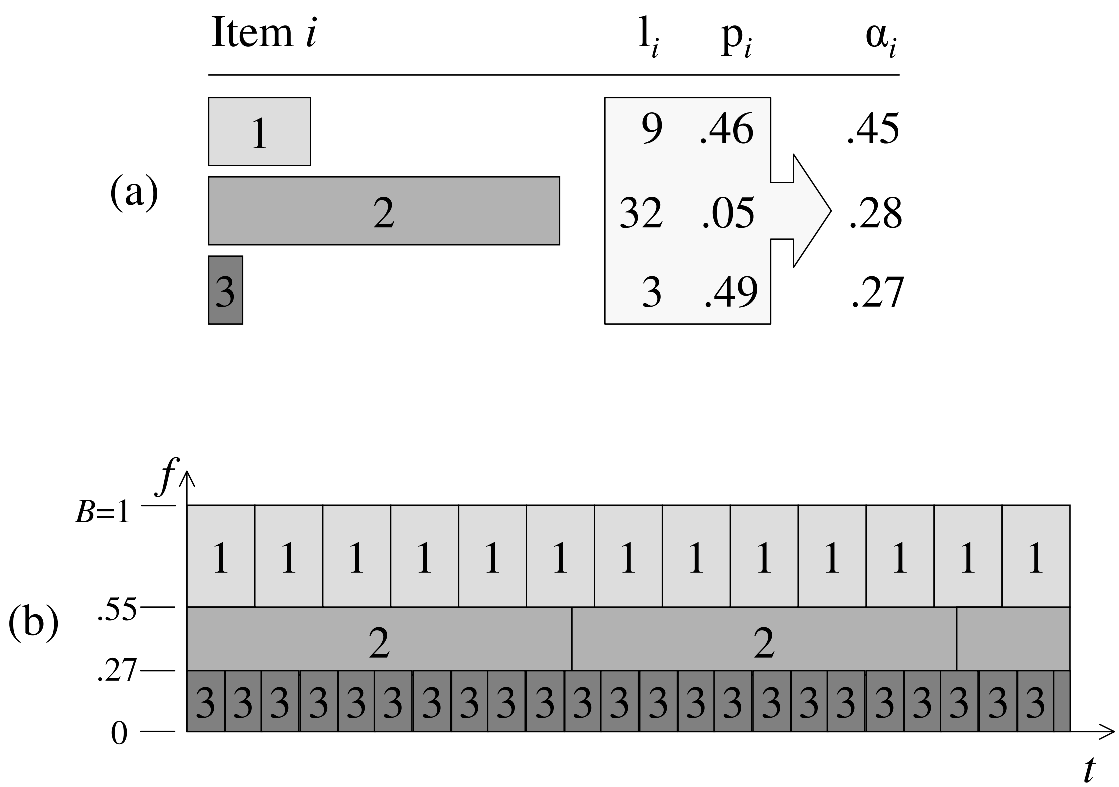

Figure 5: Optimal frequency-division scheduling: (a) computing α_i's for the items (b) the resulting schedule.

Case 1 $\left(0 < \alpha_1 \leq \frac{1}{2}\right)$

When all α_i's are at most $\frac{1}{2}$, we see that for the optimal frequency-division schedule, S_{FD}, $EWT_i\left(S_{FD}\right) = \frac{3}{2\alpha_i} - 1 = \left(\frac{3}{2} - \alpha_i\right) \cdot s_i \geq s_i$. Let $s'_i = 2^{\lceil log_2 s_i \rceil}$, the next power of 2 greater than or equal to s_i. Since $\sum_{i=1}^{n} \frac{1}{s_i} = 1$ and $s'_i \geq s_i$ for $1 \leq i \leq n$, it follows that $\sum_{i=1}^{n} \frac{1}{s'_i} \leq 1$. Since $\sum_{i=1}^{n} \frac{1}{s'_i}$ is the bandwidth requirement for scheduling with spacing s'_i, and 1 is the total bandwidth available, we have enough bandwidth to schedule with these spacings between items. We start with item 1 and schedule it with spacing s'_1. We then repeat for items 2 through n. Since the spacings are increasing powers of 2, we can always schedule the items with the appropriate spacing. Each s'_i is less than twice the corresponding s_i, and the expected waiting time for item i using the time-division schedule is $\frac{1}{2}s'_i < \frac{1}{2}\left(2s_i\right) = s_i$. Thus, for this time-division schedule, S_{TD}, we get $EWT_i\left(S_{TD}\right) < EWT_i\left(S_{FD}\right)$ $\forall i, 1 \leq i \leq n$. It follows that time-division is better than frequency-division when $\alpha_1 < \frac{1}{2}$.

Example 3

To see how to schedule items when $\alpha_1 < \frac{1}{2}$, we consider an example. Let $n = 3$, $\alpha_1 = .41$, $\alpha_2 = .36$, and $\alpha_3 = .23$. Then we assign $s'_1 = 4$, $s'_2 = 4$, $s'_3 = 8$, and schedule the items as shown in Figure 7. Using the frequency-division schedule, we get $EWT_1 = 2.66$, $EWT_2 = 3.17$, and $EWT_3 = 5.52$. with the time-division schedule, we get $EWT_1 = 1.4$, $EWT_2 = 1.4$, and $EWT_3 = 2.5$. Since the time-division schedule gives lower expected waiting times for each item, it follows that the time-division schedule has a lower expected waiting time for any values of the demand probabilities p_1, p_2, and p_3.

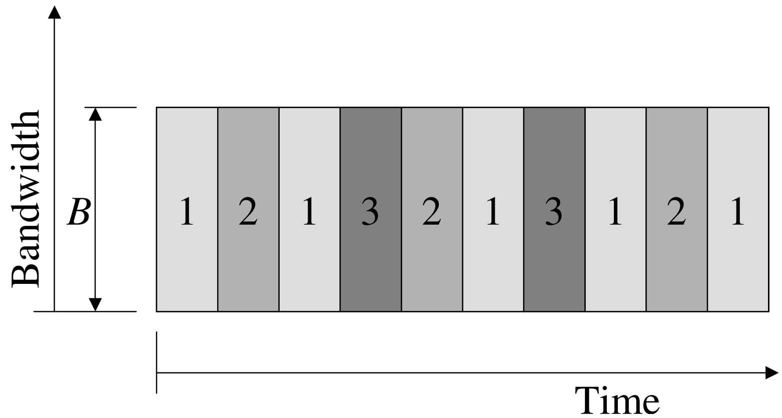

Figure 6: Time-division scheduling

Case 2 $\left(\frac{1}{2} < \alpha_1 \leq \frac{3}{4}\right)$

For the frequency-division schedule, S_{FD}, with $\frac{1}{2} < \alpha_1 \leq \frac{3}{4}$, we get $EWT_1\left(S_{FD}\right) = \frac{3}{2}s_1 - \alpha_1 s_1 = \frac{3}{2\alpha_1} - 1 \geq 1$. For our time-division schedule, S_{TD}, we let $s_1' = 2$. This gives $EWT\left(S_{TD}\right) = 1$, so $EWT\left(S_{TD}\right) \leq EWT\left(S_{FD}\right)$. For the other items, the sum of the α_i's is at most $\frac{1}{2}$, and we have $\frac{1}{2}$ the schedule to fit them in. So, we proceed as in Case 1 and round each s_i up to the next power of 2). We schedule these items as in Case 1. As before, $EWT_i\left(S_{TD}\right) < EWT_i\left(S_{FD}\right)\ \forall i, 1 \leq i \leq n$. So, time-division is better than frequency-division when $\frac{1}{2} < \alpha_1 \leq \frac{3}{4}$.

Example 4

To see how to schedule items when $\frac{1}{2} < \alpha_1 < \frac{3}{4}$, we consider an example. Let $n = 3$, $\alpha_1 = .61$, $\alpha_2 = .23$, and $\alpha_3 = .16$. Then we assign $s_1' = 2$, $s_2' = 8$, $s_3' = 8$, and schedule the items as shown in Figure 8. Using the frequency-division schedule, we get $EWT_1 = 1.46$, $EWT_2 = 5.52$, and $EWT_3 = 8.38$. with the time-division schedule, we get $EWT_1 = .83$, $EWT_2 = 3$, and $EWT_3 = 3$. Since the time-division schedule gives lower expected waiting times for each item, it follows that the time-division schedule has a lower expected waiting time for any values of the demand probabilities p_1, p_2, and p_3.

Case 3 $\left(\alpha_1 > \frac{3}{4}\right)$

Consider the time-division schedule, S_{TD}, where we broadcast item 1 r times consecutively and then leave an empty spot. This has $EWT_1\left(S_{TD}\right) = \frac{r+3}{2(r+1)}$. For the time-division schedule to have a lower EWT_1 than the frequency-division schedule, we need $\frac{r+3}{2(r+1)} < \frac{3}{2\alpha_1} - 1 \Longrightarrow r > \frac{2}{3} \cdot \frac{\alpha_1}{1-\alpha_1} - 1$, so we let $r = \left\lfloor \frac{2}{3} \cdot \frac{\alpha_1}{1-\alpha_1} \right\rfloor$. This uses $\frac{r}{r+1}$ of the schedule for item 1. This is $\frac{\left\lfloor \frac{2}{3}\cdot\frac{\alpha_1}{1-\alpha_1}\right\rfloor}{\left\lfloor \frac{2}{3}\cdot\frac{\alpha_1}{1-\alpha_1}\right\rfloor+1} < \frac{\frac{2}{3}\cdot\frac{\alpha_1}{1-\alpha_1}}{\frac{2}{3}\cdot\frac{\alpha_1}{1-\alpha_1}+1} = \frac{2\alpha_1}{2\alpha_1+3-3\alpha_1} = \frac{2\alpha_1}{2+(1-\alpha_1)} < \frac{2\alpha_1}{2} = \alpha_1$. So, we have used less than α_1 of the schedule for item 1. We let $s_i' = (r+1)\,2^{log_2 \frac{s_i}{r+1}}$, the next value larger than or equal to s_i of the form $(r+1)\,2^m, m \in \mathbb{Z}$. If $\alpha_1' = \frac{1}{s_i'}$, then $\sum_{i=2}^{n} \alpha_i' < \sum_{i=1}^{n} \alpha_i$, so our bandwidth requirement is reduced for time-division. We schedule items 2 through n as before, doubling the schedules as necessary. We have $s_i' < 2s_i \Longrightarrow EWT_i\left(S_{TD}\right) < EWT_i\left(S_{FD}\right)$. So, time-division is better than frequency-division when $\alpha_1 > \frac{3}{4}$.

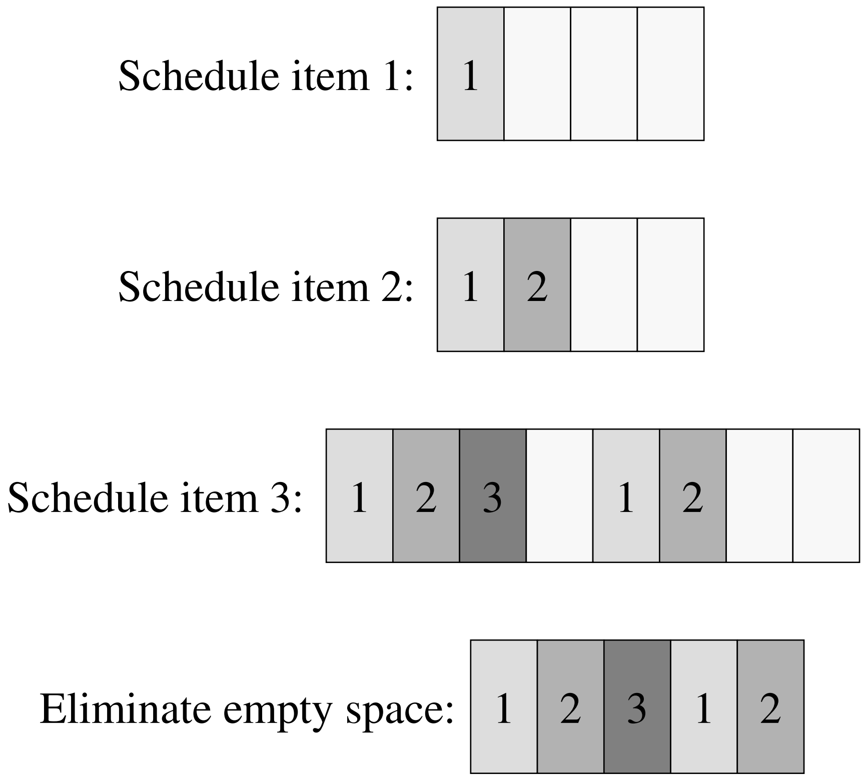

Figure 7: Generating a time-division schedule better than the frequency-division schedule with $\alpha_1 = .41$, $\alpha_2 = .36$, and $\alpha_3 = .23$.

Example 5

To see how to schedule items when $\alpha_1 > \frac{3}{4}$, we consider an example. Let $n = 3$, $\alpha_1 = .88$, $\alpha_2 = .07$, and $\alpha_3 = .05$. Then we group item 1 in blocks of 4, let $s'_2 = 20$ and $s'_3 = 20$, and schedule the items as shown in Figure 9. Using the frequency-division schedule, we get $EWT_1 = .70$, $EWT_2 = 20.4$, and $EWT_3 = 29.0$. with the time-division schedule, we get $EWT_1 = .61$, $EWT_2 = 9$, and $EWT_3 = 9$. Since the time-division schedule gives lower expected waiting times for each item, it follows that the time-division schedule has a lower expected waiting time for any values of the demand probabilities p_1, p_2, and p_3.

We see that for any value of α_1 and any frequency-division schedule we can find a time-division schedule with a lower expected waiting time. Since all the EWT_i's are lower for time-division, it follows that the expected waiting time (EWT) for the time-division schedule is lower for any values of the demand probabilities, and the theorem is proved.

5 Conclusions and Extensions

We have found the optimal bandwidth allocation for frequency-division broadcast scheduling. We have also shown a way to generate a better time-division schedule than any frequency-division schedule when the data items have equal lengths. These results are for dynamic data, data that changes from one transmission to the next.

We would like to look at what happens when the data is static, constant from one transmission to the next. We have some results for static data, but not as nice as for dynamic data. We have

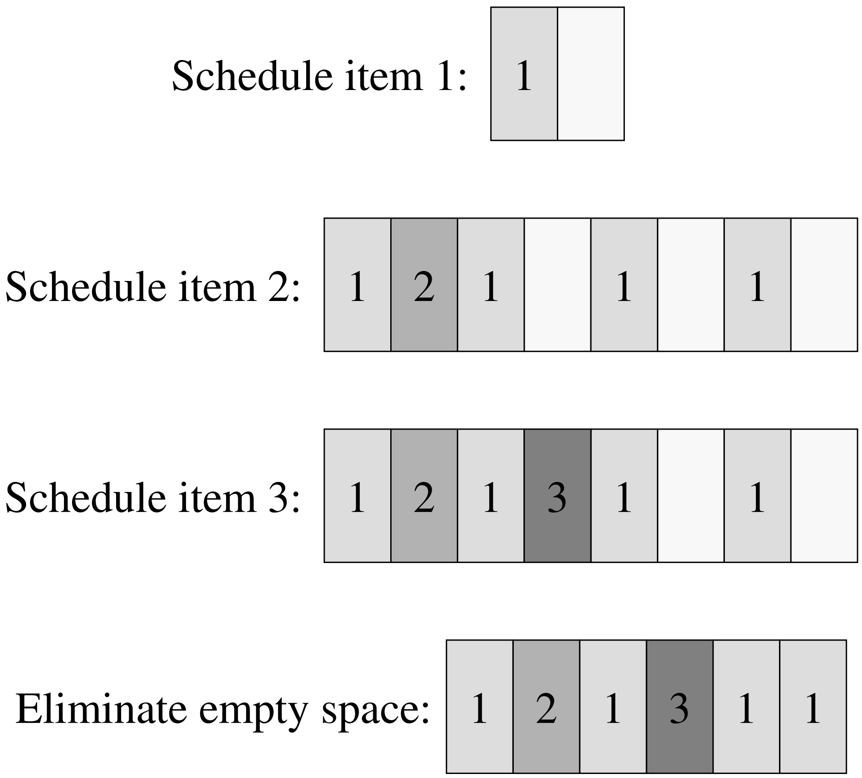

Figure 8: Generating a time-division schedule better than the frequency-division schedule with $\alpha_1 = .61$, $\alpha_2 = .23$, and $\alpha_3 = .16$.

shown that for $n = 2$ or $n = 3$ items, time-division is better than frequency-division. We would like to extend this result to arbitrary n.

We would like to extend the results of this paper to items of arbitrary length. Using time-division schedules where we send each item from start to finish with no interruptions, this is not possible. For example, consider two items of lengths $l_1 = 1$ and $l_2 = 19$, and demands $p_1 = .95$ and $p_2 = .05$. The SRR-FD schedule with $\alpha_1 = \alpha_2 = \frac{1}{2}$ is better than the optimal time-division schedule of alternately sending item 2 followed by 66 instances of item 1. However, if we split up longer items, we may be able to show that time-division is always better than frequency-division for arbitrary length items. We would also like to extend our results in [4] to arbitrary lengths and more than two items. It would also be good to find time-division schedules that not only perform better than frequency-division schedules, but are provably optimal among all time-division schedules.

References

[1] D. Aksoy, M. J. Franklin, "Scheduling for Large-Scale On-Demand Data Broadcasting", IEEE INFOCOM, San Francisco, CA, March 1998.

[2] A. Bar-Noy, R. Bhatia, J. Naor, B. Schieber, "Minimizing Service and Operations Cost of Periodic Scheduling", 9th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 1998.

[3] A. Bestavros, "AIDA-based Real-Time Fault-Tolerant Broadcast Disks", Technical Report 96-001, Computer Science, Boston University, January 5, 1996.

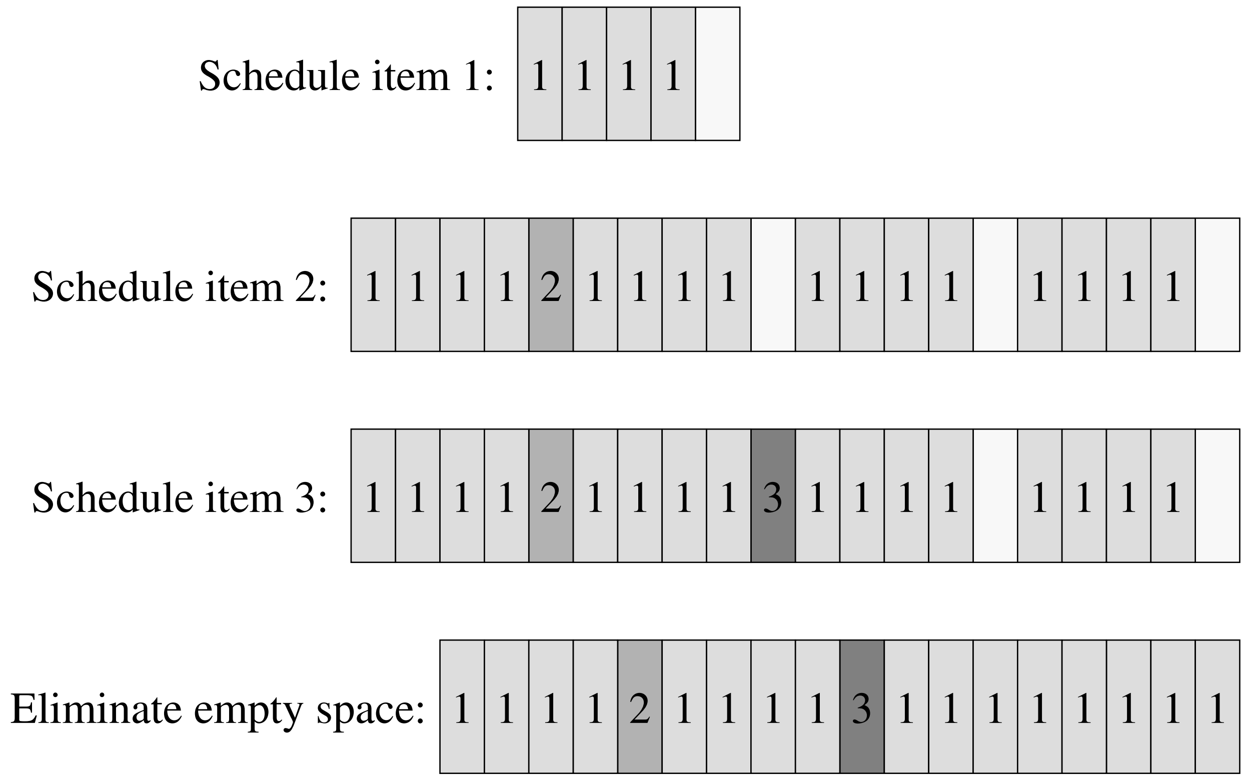

Figure 9: Generating a time-division schedule better than the frequency-division schedule with $\alpha_1 = .88$, $\alpha_2 = .07$, and $\alpha_3 = .05$.

[4] K. Foltz, J. Bruck, "Splitting Schedules for Internet Broadcast Communication", To appear in *IEEE Transactions on Information Theory*, February 2002.

[5] M. J. Franklin, S. Zdonik, "Dissemination-Based Information Systems", IEEE Data Engineering Bulletin, Vol 19, No 3, September 1996.

[6] S. Hameed, *Scheduling Information Broadcast in Asymmetric Environment*, Masters of Science thesis, Texas A&M University, May 1997.

[7] S. Hameed, N. H. Vaidya, "Log-time Algorithms for Scheduling Single and Multiple Channel Data Broadcast", MOBICOM'97, Budapest, September 1997.

[8] S. Jiang, N. H. Vaidya, "Scheduling Algorithms for a Data Broadcast System: Minimizing Variance of the Response Time", Technical Report 98-005, Computer Science, Texas A&M University, February 4, 1998.

[9] C. Kenyon, N. Schabanel, "The Data Broadcast Problem with Non-Uniform Transmission Times", Proc. of the 10th ACM-SIAM Symp. on Discrete Algorithms (SODA '99), pp. 547-556, January 1999.

[10] H. V. Leong, A. Si, "Database Caching Over the Air-Storage", *The Computer Journal* 40(7): 401-415, 1997.

[11] N. H. Vaidya, S. Hameed, "Data Broadcast in Asymmetric Wireless Environments", First International Workshop on Satellite-based Information Services (WOSBIS), Rye, NY, November 1996.

[12] S. Zdonik, M. J. Franklin, R. Alonso, S. Acharya, "Are 'Disks in the Air' Just 'Pie in the Sky'?", IEEE Workshop on Mobile Computing Systems and Applications, Santa Cruz, CA, December 1994.