

Seamless Workload Adaptive Broadcast ¹

Yang Guo, Lixin Gao[†], Don Towsley, and Subhabrata Sen[‡]

Computer Science Department University of Massachusetts Amherst MA 01003 {yguo, towsley}@cs.umass.edu	[†] ECE Department University of Massachusetts Amherst MA 01003 lgao@ecs.umass.edu	[‡] Networking Research AT&T Labs-Research Florham Park, NJ 07932 sen@research.att.com
--	--	--

Abstract

The high-bandwidth requirements and long-lived characteristics of digital video make transmission bandwidth usage a key limiting factor in the widespread streaming of such content over the Internet. A challenging problem is to develop bandwidth-efficient techniques for delivering popular videos to a large, asynchronous client population with time-varying demand characteristics. In this paper we propose *seamless workload adaptive broadcast* to address the above issues. A key component of our scheme is Fibonacci Periodic Broadcast (FPB). By introducing a feedback control loop into FPB, and enhancing FPB using techniques such as parsimonious transmission, seamless workload adaptive broadcast provides instantaneous or near-instantaneous playback services and can seamlessly adapt to workload changes. Furthermore, FPB, as proposed in this paper, is bandwidth efficient and exhibits the periodic seamless channel transition property.

I. INTRODUCTION

The high-bandwidth requirements and long-lived characteristics of digital video make transmission bandwidth usage a key limiting factor in the widespread streaming of such content over the Internet. For high-demand content, a large number of clients asynchronously issue requests to receive their chosen media streams. In addition, the demand for a particular video can vary over time, due to time-of-day (week) effects, changing popularity, etc. A challenging problem is to develop bandwidth-efficient techniques for delivering popular videos to such a large, asynchronous client population exhibiting time-varying demand characteristics. In this paper we report on the design and evaluation of such a delivery scheme.

Various techniques have been developed to reduce server and network bandwidth associated with delivering a popular video to asynchronous clients, by allowing multiple clients to receive all, or part of, a single transmission. Periodic Broadcast (PB) schemes [1–5] divide a video object into multiple segments, and continuously broadcast the segments on a set of transmission channels. Using a constant number of channels, PB can provide streaming video with a pre-determined playback startup delay to an arbitrary number of clients. Other proposed techniques, such as patching and stream merging [6–13], are not as bandwidth efficient as PB when the client arrival rate is high.

Research on Periodic Broadcast has focused on further improving its efficiency – to reduce the server bandwidth requirement with a pre-determined playback delay while keeping the client side resource requirement, such as clients’ receiving bandwidth or work-ahead buffer size, low. However the Periodic Broadcast schemes proposed so far exhibit the following drawbacks:

- *Workload insensitivity.* A Periodic Broadcast scheme is essentially an open loop scheme that does not adapt to changing workload demand. PB transmits all segments and uses the same amount of bandwidth

¹This research was supported in part by the National Science Foundation under NSF grants EIA-0080119, NSF ANI-9973092, ANI9977635, ANI-9977555, and ANI9875513. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

regardless of the demand for the video. PB is designed to serve popular videos. However in reality video popularity changes over time. Furthermore, the popularity of videos often cannot be determined in advance.

- *Delayed playback.* Clients in Periodic Broadcast experience a playback delay, which can be significant if the number of channels used is small.

A broadcast scheme that can adapt to the dynamically changing workload and offer instantaneous, or near-instantaneous playback is desirable. In this paper, we propose *seamless workload adaptive broadcast* to address above issues. Seamless workload adaptive broadcast is based on the *workload adaptive broadcast architecture* and *Fibonacci PB (FPB)*. Our workload adaptive broadcast architecture is centered around a PB scheme, with the addition of following techniques.

- *Parsimonious transmission.* The server transmits a segment only if it is required by some clients.
- *Workload adaptation.* Addition of a control loop into Periodic Broadcast helps Periodic Broadcast adapt to the workload. The server collects client arrival information, and then dynamically adjusts the number of channels used in PB to minimize the overall bandwidth usage.
- *Instantaneous playback.* This technique enables the instantaneous or near-instantaneous playback in workload adaptive broadcast.

We introduce the Fibonacci Period Broadcast (FPB) scheme, which is especially suitable for workload adaptive broadcast architecture. The sizes of the video segments delivered over the channels are proportional to the Fibonacci number series, hence the name. A PB scheme exhibits the seamless transition property if it can periodically change the number of channels without disrupting existing clients' reception and without any additional channels. FPB exhibits the *seamless transition property* and is as bandwidth efficient as other similar PB schemes. The seamless transition property provides the opportunity to adjust the number of server channels smoothly. Thus FPB is especially suitable for workload adaptive broadcast. We also derive a recursive formula to calculate the average bandwidth usage, which can be used to determine when to add or remove server channels.

Simulation studies show that FPB is more bandwidth efficient than other PB schemes with the same client-side network bandwidth requirement. Parsimonious FPB saves network bandwidth when the client request rate is low. Finally, we show that the seamless workload-adaptive scheme serves a video whose popularity changes. The bandwidth usage in seamless workload-adaptive broadcast is proportional to the workload.

The remainder of the paper is organized as follows. In Section 2 we describe the architecture of workload adaptive broadcast. In Section 3 we present the seamless workload-adaptive scheme. Section 4 is dedicated to performance evaluation. Conclusions and future work are included in Section 5.

II. WORKLOAD ADAPTIVE BROADCAST ARCHITECTURE

In this section, we describe the workload adaptive broadcast architecture. The workload adaptive broadcast architecture is centered around a Periodic Broadcast scheme coupled with additional features such as parsimonious transmission, dynamic channel adjustment, and instantaneous or near-instantaneous playback. These features enable it to perform well even under changing workloads.

Fig. 1 depicts the architecture of the workload adaptive broadcast. The server consists of two components: a modified PB scheduler and a workload adaptor. Below we describe them respectively.

- **Modified PB scheduler.** This component serves the video using a modified PB scheme that (1) provide instantaneous playback service; and (2) saves the network bandwidth by only transmitting segments that are needed by clients (parsimonious transmission).

To illustrate how instantaneous playback is achieved, suppose a video clip is divided into equal size




Fig. 1. Workload adaptive broadcast architecture

segments. Fig. 2 illustrates an example where the video is divided into 11 equal size segments. When a client arrives, segment A is unicast to the client immediately to enable instantaneous playback. The client then starts to receive the data from the modified PB scheme, which uses one more channel for segment B . Since segments A and B are of the same size, the client can start to receive segment B while playing back segment A . Moreover, since segment B and segment 1 are of the same size, they can be received sequentially. The above device helps to achieve the instantaneous playback while the client still listens to the same number of server channels as before. For instance, in Fig. 2, the client 1 receives segment A immediately, and then receives shadowed segment B , and backslashed segment 1, etc. from multicast channels.

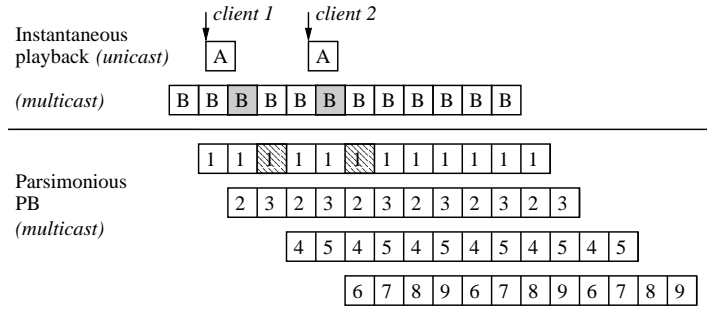


Fig. 2. Modified PB scheme

Considering the network's propagation delay and client side's buffering delay, a short delay on the order of seconds can be treated as near-instantaneous playback. If a small playback delay is allowed, the component can batch the requests and multicast segment A , which can further reduce the bandwidth usage. In the following discussion, we always assume that instantaneous playback is required. The workload adaptive broadcast architecture applies to both instantaneous and near-instantaneous playback schemes.

• **Workload adaptor.** The workload adaptor collects the client arrival rate information, and determines the number of channels required by the modified PB scheduler in order to minimize overall bandwidth usage.

We use an exponential smoothing algorithm to estimate the average client request rate. The number of arrivals is periodically collected. Denote by $\hat{\lambda}_n$ as the arrival rate after the n -th update period, then

$$\hat{\lambda}_{n+1} = \omega \hat{\lambda}_n + (1 - \omega) A_n / \Delta. \quad (1)$$

where A_n is the number of arrivals during the n -th period and Δ is the period length. The weight, ω , and the update rate, $1/\Delta$, determine how quickly the average arrival rate converges to the current arrival rate.

The number of channels used by the modified PB scheduler, denoted by K , is determined by the average arrival rate. Let $B(K, L, \lambda)$ be the average bandwidth usage of workload adaptive broadcast where K is the

number of channels allocated and L is the length of the video clip. The workload adaptor chooses the number of channels, K^* , so as to minimize the overall average bandwidth usage, i.e., $K^* = \operatorname{argmin}_K B(K, L, \lambda)$. If a change in number of channels is necessary, the adaptor notifies the modified PB scheduler to make the change. During the transition period, it is desirable that clients not be disrupted and the service of other videos not be affected. In the next section, we will introduce the FPB scheme, which provides the property of smooth channel transition.

III. SEAMLESS WORKLOAD ADAPTIVE BROADCAST

In this section we first present FPB, an efficient PB scheme that exhibits the seamless transition property. We then describe seamless workload adaptive broadcast using the FPB.

A. Fibonacci Periodic Broadcast (FPB)

The FPB scheme has the segmentation series of $1, 2, 3, 5, 8, \dots$. This corresponds to the Fibonacci series with the first two numbers, 0 and 1, excluded. Channel n , $n = 1, 2, 3, \dots$, is responsible for delivering consecutive F_n segments to clients, where F_n is defined as

$$\begin{cases} F_0 = F_1 = 1 \\ F_{n+2} = F_{n+1} + F_n, \quad (n = 1, 2, 3, \dots) \end{cases} \quad (2)$$

Below we describe the server's transmission schedule and client's reception schedule respectively.

A.1 Server's broadcasting schedule

Suppose the FPB scheme uses K channels to transmit a video clip of length L . The n -th channel, $1 \leq n \leq K$, is responsible for delivering F_n segments to clients, from segment $\sum_{i=1}^{n-1} F_i + 1$ to segment $\sum_{i=1}^n F_i$. We use $[\sum_{i=1}^{n-1} F_i + 1, \sum_{i=1}^n F_i]$ to represent these consecutive segments. The FPB scheme consists of a start rule, a repeat rule, and a transmission pattern within a period.

Start rule. The transmission of channel 1 starts first. The n -th channel starts transmission after the $(n - 1)$ -th channel completes the transmission of F_{n-1} segments.

Repeat rule. Each channel repeats its transmission pattern once every F_K segments. We call F_K the *period* of the FPB scheme.

Transmission pattern within a period. The first channel repeatedly sends out the first segment F_K times. For channel n , $n = 2, 3, \dots, K - 1$, the transmission pattern is illustrated in Fig. 3. Channel n first transmits F_n segments, segment $[\sum_{i=1}^{n-1} F_i + 1, \sum_{i=1}^n F_i]$. It then transmits $K - n$ batches of segments, where batch i , $i = 1, 2, \dots, K - n$, consists of F_{K-2+i} segments. Since $F_n + \sum_{i=n-1}^{K-2} F_i = F_K$, the period is F_K . Each batch repeats the segments that have been transmitted from the beginning of this period. Finally, the last channel, channel K , sends out the last F_K segments of the clip sequentially.




Fig. 3. Transmitting Pattern of Channel n (one period)

Figure 4 gives an example of FPB using six channels. The video clip is divided into $\sum_{i=1}^6 F_i = 32$ segments. The period is $F_6 = 13$ segments. The transmission pattern is as described above. For instance,

the third channel is responsible for transmitting segment $\{4, 5, 6\}$ to clients. It starts by sending out segment $\{4, 5, 6\}$, and followed by three batches, $\{4, 5\}$, $\{4, 5, 6\}$, and $\{4, 5, 6, 4, 5\}$, respectively. Each batch repeats the segments that has been transmitted from the beginning of this period. We call the collection of one period of K channels a K -channel cluster of FPB scheme. All K -channel clusters are identical and independent of each other. In fact clients that start to receive segments within a cluster only fetch the data from the same cluster. Therefore it suffices to describe the client's reception schedule in one cluster.

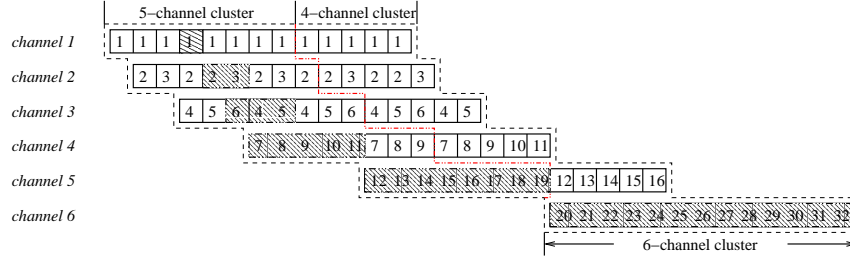


Fig. 4. A 6-channel Cluster and Its Sub-clusters in FPB

A.2 Client's reception schedule

The cluster exhibits a recursive structure. For instance, the 6-channel cluster in Fig. 4 consists of a 5-channel cluster and a 4-channel cluster. The 5-channel and 4-channels are further subdivided into clusters. We explore the cluster's recursive structure in the FPB scheme, and present an algorithm that generates the client's reception schedule.

Reception schedule in 1-channel cluster. This is a trivial case. The client receives the first segment immediately.

Reception schedule in 2-channel cluster. Denote by T the starting time of the cluster, and by P the arrival time of the client. We use the segment length as the time unit. All clients arriving during a segment will be batched and served together at the starting time of the next segment. Hence we use the starting time of the next segment as the arrival time of these clients.




Fig. 5. A 2-channel Cluster

If $P = T$, clients receive the first instance of segment 1, and continue to receive segment 2 and 3 from channel 2. If $P = T + 1$, clients receive the second instance of segment 1, and simultaneously receives segment 2. Segment 3 is received after segment 2. In both cases, clients listen to at most two channels simultaneously.

Reception schedule in K -channel cluster ($K > 2$). Above we have shown that there is a valid reception schedule for a 1-channel and 2-channel cluster where clients listen to at most 2 channels simultaneously. Suppose there is a valid reception schedule for a $(K - 1)$ -channel cluster and a $(K - 2)$ -channel cluster. In the following we show that there exists a valid reception schedule for a K -channel cluster. By induction, a valid reception schedule exists for an arbitrary cluster.

A K -channel cluster has two sub-clusters, a $(K - 1)$ -channel cluster and a $(K - 2)$ -channel cluster. If a client arrives during the first F_{K-1} segments, it receives the segments associated with the first $K - 1$ channels according to the reception schedule for the $(K - 1)$ -channel cluster. Once these have been received, the client receives the F_K segments associated with the K -th channel. Since the reception of segments from the K -th channel occurs after the reception from the first $K - 1$ channels finishes, clients listen to at most 2 channels.

If a client arrives during the F_{K-2} segments, it receives the segments associated with the $K - 2$ channels according to the reception schedule for the $(K - 2)$ -channel cluster. Once these have been received, the client tries to receive segments from $(K - 1)$ -th and K -th channel, without violating the “listening to at most 2 channels” rule. The client can start to listen to the $(K - 1)$ -th channel once the $(K - 3)$ -th channel finishes the transmission at time $T + F_{K-2} - 2 + F_K$, where T is the starting time of the cluster. It can be shown that clients must receive all F_{K-1} segments from channel $K - 1$. The completion time of the $(K - 2)$ -th channel coincides with the starting time of the K -th channel; thus the clients are able to fetch the last F_K segments from the K -th channel.

Pseudo-code for generating the client reception schedule is included in [14]. As an example, suppose the client starts at the time of the 4th segment in channel 1 (backslashed segment in Fig. 4). Since it falls into a 5-channel cluster, the client receives 13 segments, 20 – 32, from channel 6. Within the 5-channel cluster, the 4th segment belongs to the 4-channel cluster, thus it receives 8 segments, 12 – 19, from channel 5. Within this 4-channel cluster, the 4th segment belongs to the later 2-channel cluster, instead of the leading 3-channel cluster. Thus it receives 5 segments, 7 – 11, from channel 4, and 3 segments from channel 3, in the order of segment 6, 4, and 5. Since the 4th segment is the first segment in a 2-channel cluster, the client obtains the first segment immediately, and segments 2 and 3 in the following slot. The segments received by this client are marked as backslashed segments in Fig. 4. Note that if clients start the reception from the first segment in a cluster, they can receive the entire video listening to one channel at a time and no client-side buffer is required.

A.3 Seamless transition property

FPB exhibits the seamless channel transition property. Assume that it uses a fixed number, say K , of channels, and the newly assigned number of channels is K' . During the channel transition period, we require that (1) the clients already starting their service not experience any disruption during the transition; (2) the newly arrived clients use the FPB scheme with K' channels; (3) the total number of channels used during the transition period is no larger than $\max\{K, K'\}$. We call a transition satisfying the above conditions a *seamless transition*.

A naive solution is to allocate another set of K' channels for newly arrived clients. The previous K channels are held until all old clients are served. The solution requires $K + K'$ channels and wastes the bandwidth during the transition period. Moreover, if the server supports multiple video clips, the channel transition can lead to a resource deadlock problem. We state the following result with the proof included in [14].

Theorem 1: A seamless transition can be achieved at a cluster boundary in FPB scheme.

B. Seamless Workload Adaptive Broadcast

Seamless workload adaptive broadcast uses parsimonious FPB scheme in the modified PB scheduler. The channel transition is made at the boundary of each cluster when necessary. Below we first describe how to calculate the average bandwidth usage in seamless workload adaptive broadcast. This is used to create a

table of bandwidth usages indexed by the normalized workload, the product of video request rate and video length. The need for a channel transition is determined from a table lookup. If necessary, the workload adaptor performs a transition at the boundary of the cluster, leading to a smooth transition.

Since two extra segments are needed to provide instantaneous playback (see Fig. 2), the segment size in seamless workload-adaptive broadcast, $g'(K, L)$, is

$$g'(K, L) = \frac{L}{\sum_{i=1}^K F_i + 2}. \quad (3)$$

On average there are $\lambda g'(K, L)$ arrivals during a segment. For each arrival, the modified PB scheduler transmits segment A, and, hence, the average number of transmitted segments from the modified PB scheduler in a period is $\lambda g'(K, L) F_K$. Denote by p the probability of an arrival in a segment. Suppose the client arrival process is Poisson with arrival rate λ . We have $p = 1 - e^{-\lambda g'(K, L)}$ for all segments because of the memoryless property of Poisson process. For the PB component, an average number of $p F_K$ copies of segment B are transmitted. Denote by $C(K, p)$ the average number of segments transmitted during a K -channel cluster by Parsimonious FPB. The average total number of segments transmitted is $\lambda F_K g'(K, L) + C(K, p) + p F_K$. Hence the average number of busy channels, $B(K, \lambda, L)$, is

$$B(K, L, \lambda) = \lambda g'(K, L) + p + \frac{C(K, p)}{F_K}. \quad (4)$$

We would like to choose a value of K , K^* , that minimizes the average number of busy channels, i.e., $K^* = \operatorname{argmin} B(K, L, \lambda)$.

Theorem 2: In parsimonious FPB, the average number of segments transmitted in a K -channel cluster, $C(K, p)$, satisfies the recursion

$$\begin{aligned} C(K, p) = & C(K-1, p) + C(K-2, p) + (1-p)^{F_{K-1}}(1 - (1-p)^{F_{K-2}})F_{K-3} \\ & + (1 - (1-p)^{F_{K-2}})F_{K-2} + (1 - (1-p)^{F_K})F_K \end{aligned} \quad (5)$$

for $K \geq 3$, where $C(1, p) = p$, and $C(2, p) = 6p - 2p^2$.

The proof is included in [14].

No closed-form solution of K^* exists. Fig. 6 shows the average number of busy channels as a function of number of channels, K , for different client arrival rates. As the arrival rate increases and the number of channels increases, the average number of busy channels decreases substantially. However there exists a number of channels above which the improvement rapidly decreases.

Fig. 7 plots the optimal number of channels as the normalized workload varies. The optimal number of channels is defined as the least number of channels such that the bandwidth usage is within 1% of the minimum bandwidth usage. Here the normalized workload is the product of the client request arrival rate and the video length. The curve exhibits a staircase shape. We can determine the range of normalized workloads over which the same number of channels are required and represent it in a table. The workload adaptor can choose the number of server channels from the table to reduce bandwidth usage while maintaining its own schedule as well as clients' reception schedule as simple as possible. If a certain playback delay can be treated as near-instantaneous playback, the modified PB scheduler can batch the requests and multicast the segment A. In the following section, however, we assume instantaneous playback is desired. We expect a similar result when small playback delay is allowed.




Fig. 6. Average Number of Busy Channels vs. Number of Channels with Different Client Workload

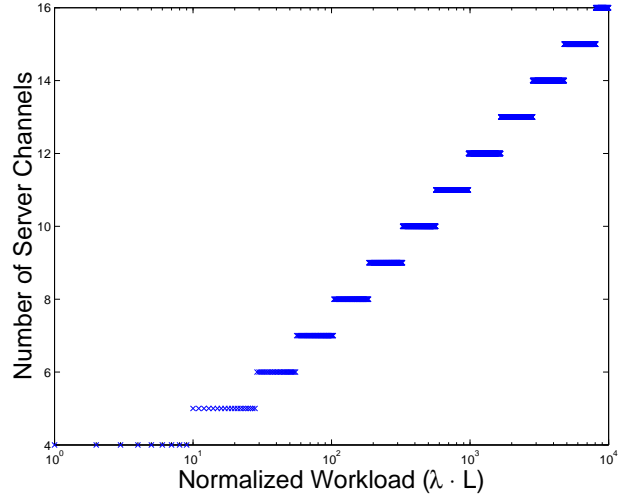


Fig. 7. Optimal Number of Server Channels vs. Normalized Workload

IV. PERFORMANCE EVALUATION

We evaluate the seamless workload-adaptive broadcast scheme from the following three perspectives: (1) how FPB compares with other PB schemes; (2) whether and how much parsimonious FPB can save server bandwidth when the client request rate is low; and (3) how seamless workload adaptive broadcast adapts to changing video popularity. We show that FPB is more efficient than other popular PB schemes with the same client side network bandwidth requirement. Parsimonious FPB scheme uses less bandwidth when the client request rate is low. To evaluate the seamless workload adaptive broadcast, we use a workload whose rate changes dramatically throughout a day. Simulation results show that seamless workload adaptive broadcast adapts nicely to the changing workload.

- **Comparison of PB schemes.** Fig. 8 illustrates the server bandwidth requirement (number of channels) vs. the startup delay represented in fraction of video length. We compare FPB with dynamic skyscraper, skyscraper, and GDB3, which all require clients listen to 2 channels simultaneously. For the same amount of playback delay, FPB uses less bandwidth than other schemes.

- **Efficiency of Parsimonious FPB.** Fig. 9 shows the average number of busy channels vs. the client request rate. Three curves corresponds to the cases where 7, 10, and 20 server channels are used, respectively. As the client request rate increases, the average number of busy channels also increases. The reason behind this is that more segments are transmitted as more requests arrive. Eventually, all segments need to be sent out, and all channels assigned to the PB scheme are used. Thus the number of busy channels reaches the number of channels assigned to the PB scheme.

- **Performance of seamless workload adaptive broadcast.** Finally we investigate the performance of the seamless workload adaptive scheme. Fig. 10 depicts the client arrival rate during 24-hour period. The arrival process is Poisson with a time varying rate. During peak hours (from 10am to 4pm), the rate is around 15 requests/min, while during off-peak hours, the rate is around 0.3 requests/min. The dotted line is the estimated client arrival rate from the workload estimator. The exponential smoothing average algorithm (formula (1)) is used to keep track of the client arrival rate. In this experiment, w is set to be 0.1. The average client arrival rate is updated once every minute. We can see that the workload estimator does a nice job keeping track of the actual arrival rate, filtering out the short-term rate change.

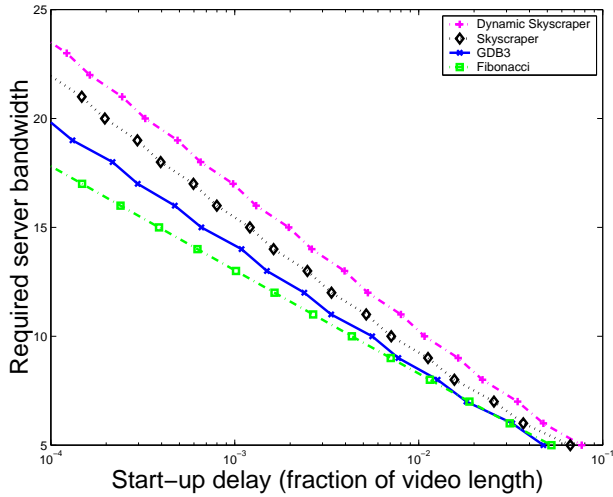


Fig. 8. Comparison of PB schemes (requiring clients listen to 2 channels simultaneously)

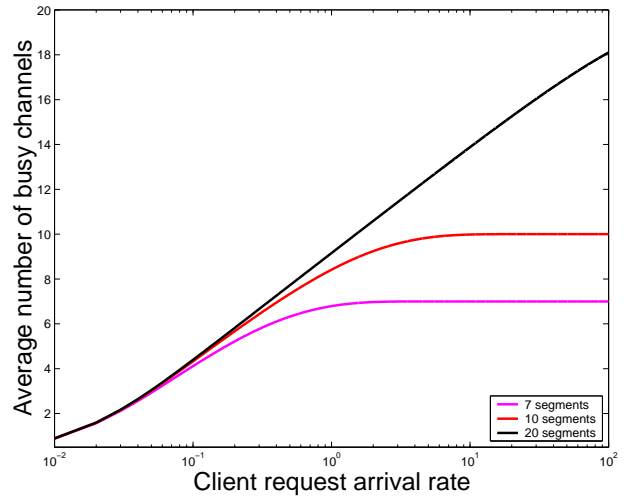


Fig. 9. Average Number of Busy Channels vs. Client Arrival Rate in Parsimonious FPB

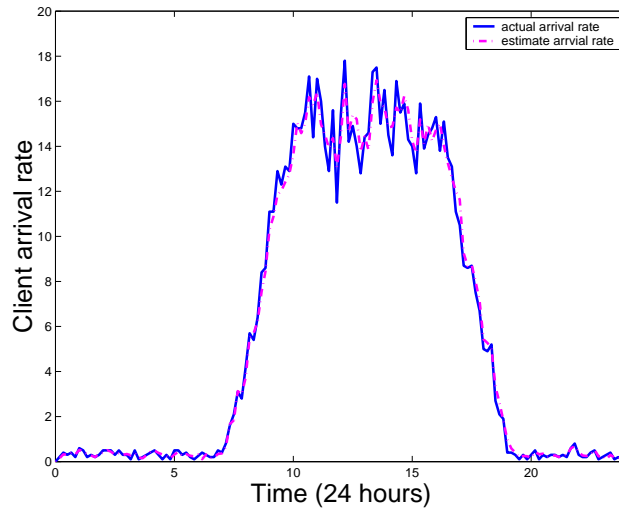


Fig. 10. Client arrival rate vs. time of the day

The seamless workload adaptive broadcast chooses the number of channels used in FPB based on the estimated client arrival rate to minimize the required bandwidth to serve the clients. Fig. 11 shows the number of channels chosen at different times during the day. More channels are used when the client arrival rate is high. The shape of the curve resembles the client arrival rate process.

Fig. 12 depicts the number of active channels (server bandwidth usage) over time. The bandwidth usage in seamless workload adaptive broadcast is proportional to the workload. Seamless workload adaptive broadcast adapts to the workload by adjusting the number of channels used in the PB component. The dashed line in the figure is the bandwidth consumed by the PB component alone (not including the bandwidth used for instantaneous playback). The difference between the solid and dashed lines is the bandwidth required to provide instantaneous playback. There are more arrivals during the peak hour. The seamless workload adaptive broadcast increases the number of channels used in the peak hour, thus decreasing the size of a segment. The result is that the bandwidth used to enable the instantaneous playback does not increase dramatically; thus the overall bandwidth usage remains low.




Fig. 11. Number of channels used in FPB vs. time of the day




Fig. 12. Number of active streams vs. time of the day

V. RELATED WORK

Several works [15–17] in the past addressed the problem of changing workload adaption. [15] proposes to use a PB scheme to serve popular videos and dynamically change the number of channels assigned to a video based on the level of demand. The proposed PB scheme possesses the seamless channel transition property; thus the channel change can be seamlessly performed. However the scheme in [15] requires the client to be able to listen to the same number of channels as used by the server. Most importantly, if a video turns cold, the PB scheme becomes inefficient. Also a playback delay is incurred by every client.

Both [16] and [17] use variations of the *batching* technique to tackle the problem of changing workload adaption. Batching serves multiple requests with a single multicast stream. [16] introduces rate-based channel allocation scheduling into batching to account for the changing workload. Since batching doesn't allow clients to prefetch data, it is not as bandwidth efficient as PB. [17] proposes a hybrid scheme that is a combination of batching and PB. Popular videos are handled by PB, while batching is used for less popular videos. The popularity of the videos is periodically re-evaluated to determine the group of videos that are served by PB. It is challenging to determine whether a video should be served by PB or by batching, and both schemes require clients to wait for a period of time before being served. In contrast, seamless workload adaptive broadcast provides a unified approach that performs well for both popular and less popular videos, and offers instantaneously playback if necessary.

[18] and [4] propose the PB scheme that also has the Fibonacci series as segmentation series. From the point of view of server bandwidth efficiency, they are equivalent to the FPB scheme proposed in this paper. However both the server-side broadcasting and the client-side reception schedules are quite different. Neither of these PB schemes exhibit the seamless channel transition property, essential for seamless workload adaptive broadcast. Also our FPB scheme can support clients with few resources by scheduling them at the beginning of a cluster.

VI. CONCLUSIONS AND FUTURE WORK

In this paper we present a workload adaptive broadcast architecture and seamless workload adaptive broadcast based on FPB to provide VoD service to a large, asynchronous client population with time-varying

workload. By introducing the feedback control loop into the PB scheme, and enhancing the PB scheme by the techniques such as parsimonious transmission and instantaneous playback technique, the seamless workload adaptive broadcast provides instantaneous or near-instantaneous playback service and can adapt seamlessly to workload changes. Simulation experiments show that the required bandwidth is proportional to the client request rate. The FPB scheme proposed in this paper is bandwidth efficient and has the seamless channel transition property.

Future research can proceed along several avenues. Workload adaptive broadcast is a framework that can be used by many PB schemes. We would like to explore the possibility of using other PB scheme in workload adaptive broadcast architecture. Secondly, it will be interesting to compare the performance with other delivery schemes such as patching and stream merging schemes. Finally, implementation of seamless workload adaptive broadcast in the test-bed can further help us to evaluate the scheme in a practical setting.

REFERENCES

- [1] K. Hua and S. Sheu, "Skyscraper broadcasting: A new broadcasting scheme for metropolitan video-on-demand systems," in *Proc. ACM SIGCOMM*, September 1997.
- [2] L. Gao, D. Towsley, and J. Kurose, "Efficient schemes for broadcasting popular videos," in *Proc. Inter. Workshop on Network and Operating System Support for Digital Audio and Video*, July 1998.
- [3] A. Hu, "Video-on-demand broadcasting protocols: A comprehensive study," in *Proc. IEEE INFOCOM*, April 2001.
- [4] A. Mahanti, D. L. Eager, M. K. Vernon, and D. Sundaram-Stukel, "Scalable on-demand media streaming with packet loss recovery," in *Proc. SIGCOMM 2001*, August 2001.
- [5] S. Sen, L. Gao, and D. Towsley, "Frame-based periodic broadcast and fundamental resource tradeoffs," in *Proc. IEEE International Performance Computing and Communications Conference*, April 2001.
- [6] S. Carter and D. Long, "Improving video-on-demand server efficiency through stream tapping," in *Proc. International Conference on Computer Communications and Networks*, 1997.
- [7] K. Hua, Y. Cai, and S. Sheu, "Patching: A multicast technique for true video-on-demand services," in *Proc. ACM Multimedia*, September 1998.
- [8] L. Gao and D. Towsley, "Supplying instantaneous video-on-demand services using controlled multicast," in *Proc. IEEE International Conference on Multimedia Computing and Systems*, 1999.
- [9] D. Eager, M. Vernon, and J. Zahorjan, "Bandwidth skimming: A technique for cost-effective video-on-demand," in *Proc. SPIE/ACM Conference on Multimedia Computing and Networking*, January 2000.
- [10] A. Bar-Noy and R. E. Ladner, "Competitive on-line stream merging algorithm for media-on-demand," in *Proc. of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, January 2001.
- [11] E. C. Jr., P. Jelenkovic, and P. Momcilovic, "The dyadic stream merging algorithm," in *Proc. of Web Caching and Content Distribution*, June 2001.
- [12] Y. Guo, S. Sen, and D. Towsley, "Prefix caching assisted periodic broadcast for streaming popular videos," in *Proc. International Conference on Communications*, April 2002.
- [13] B. Wang, S. Sen, M. Adler, and D. Towsley, "Proxy-based distribution of streaming video over unicast/multicast connections," in *Proc. IEEE Infocom 2002*, to appear.
- [14] Y. Guo, L. Gao, D. Towsley, and S. Sen, "Seamless workload adaptive broadcast," Tech. Rep. 02-05, Department of Computer Science, University of Massachusetts Amherst, 2002.
- [15] Y. Tseng, C. Hsieh, M. Yang, W. Liao, and J. Sheu, "Data broadcasting and seamless channel transition for highly-demanded videos," in *IEEE INFOCOM 2000*, pp. 727–736, March 2000.
- [16] K. Almeroth, A. Dan, D. Sitaram, and W. Tetzlaff, "Long term resource allocation in video delivery systems," in *Proc. IEEE INFOCOM*, April 1997.
- [17] J. Oh, K. A. Hua, and K. Vu, "An adaptive video multicast scheme for varying workloads," to appear in *ACM/Springer Multimedia Systems*, 2002.
- [18] J.-F. Paris and D. Long, "Limiting the receiving bandwidth of broadcasting protocols for video-on-demand," in *Proceedings of the Euromedia 2000 Conference*, May 2000.