Tailored transmissions for efficient Near-Video-On-Demand service*

Yitzhak Birk & Ron Mondri
Technion
Haifa 32000, Israel
birk@ee.technion.ac.il mondri@tx.technion.ac.il

Abstract

Near-Video-On-Demand (NVOD) entails the provision of viewing flexibility that approaches that of VOD to an unlimited number of viewers of a "hot" movie at a fixed cost to the provider. This presents an exciting opportunity to service providers, especially when employing broadcast channels such as cable television and satellites. Unlike previous NVOD schemes, which were all highly structured, we employ an algorithmic approach to produce Tailor-Made transmission schemes that dominate all previous ones in both resource consumption and flexibility. The ability to tailor the transmissions to match resource availability is often key to the practicality of the service in a given environment. The design algorithms have been coded, and a prototype has been constructed.

Keywords: *Video-on-demand, NVOD, Video server.*

1. Introduction

Video-On-Demand (VOD) is the online version of traditional video-rental services. Each viewer receives a dedicated "copy" of the movie and can view it in a flexible manner, including the ability to pause and resume, rewind, and possibly even fast-forward. With VOD, the "rental" operation is essentially instantaneous, and viewing can begin within seconds of the decision to view.

The "copy" in VOD is a dedicated video stream. This stream is generated by a video server and sent to the viewer over a communication network. An important advantage of VOD over tape rental is the great flexibility in allocation of resources: the maximum number of concurrent viewers is independent of viewing choices, and is limited only by the server's total streaming capacity. The required bandwidth resources both in the server and in the communication network are proportional to the number of concurrent viewers.

The design of VOD systems has focused on the video servers themselves, with special attention to efficient bandwidth utilization of the storage devices. Related research pertains to data streaming over networks.

Near-Video-On-Demand (NVOD)

There are important situations in which many viewers wish to view the same content during the same period of time, but not simultaneously. One example is a newly released "hot" movie that is moreover advertised heavily.

Although VOD could in principle be used to address such situations, it should be noted that even the total (over all movies) number of concurrent viewers may be much higher than usual. It would therefore be very costly if not impossible to design the infrastructure (server and communication network) for such peaks. The challenge is to exploit the knowledge that there are many concurrent viewers of the "hot" movie in order to provide the service more efficiently. The solutions entail providing to an unlimited number of viewers of the same movie similar service flexibility to that of VOD at a reasonable cost to the server and communication network. Ideally, this cost is independent of the number of viewers. The service is dubbed "Near Video On Demand". For viewing of movies, we define "near" to mean commencement of viewing within 30 seconds of viewer request, and the ability to pause and resume at any time. Rewind and fast-forward functions are not a requirement. NVOD presents an exciting business opportunity to service providers: at long last, they will be able to collect unbounded revenues for "hot" titles while keeping costs fixed.

With NVOD, the nature of the communication network affects the extent to which the resource expenditure can be independent of the number of viewers. To this end, we distinguish between two types of networks based on the nature of the underlying physical layer: point-to-point networks, e.g., ones using the local-loop telephone lines, and broadcast networks that normally employ the cable television lines or satellites. We focus on broadcast networks. Here, the bandwidth required for the transmission of the same information to multiple recipients

---

* This work was supported in part by the Multimedia On-line Services Consortium, administered by the office of the Chief Scientist of the Israeli ministry of Industry and Commerce.

0-7695-0253-9/99 $10.00 © 1999 IEEE

can be independent of the number of recipients. The down side of broadcast networks is that the total available bandwidth is often lower than with point-to-point networks, since there is no spatial multiplexing.

NVOD schemes may be divided into two categories: open-loop schemes and closed-loop schemes. With open-loop schemes, there is no feedback from the viewing client to the server, so neither server transmissions nor routing on the network are affected by viewer actions (other than the possible effect on routing due to a viewer joining a multicast group). Open-loop schemes lend themselves most naturally to broadcast networks, and are uniquely suited to such networks that have only one-way communication. This is the common case in satellite-based information dissemination networks. Closed-loop schemes [1][2] permit some feedback that allows the server to adapt to client requests. This paper considers open-loop schemes with broadcast networks.

Previously proposed open-loop schemes

Consider an NVOD system that should provide a movie of length *L* seconds in a manner that permits viewing to commence within *D* seconds after viewer request. The straightforward approach for providing such a service at a fixed cost is to start a copy of the movie every *D* seconds. Viewing would simply entail each client's choosing one of the streams. The aggregate transmission rate (in steady state) would be *L/D* streams. This is independent of the number of viewers, but could be prohibitively expensive. For example, a 100-minute movie with a viewing delay of up to 30 seconds would require 200 streams, each at the original video rate. We refer to this scheme as the Baseline scheme. Over the past several years, much research has been devoted to NVOD schemes that require far less bandwidth while offering the same quality of service.

Recently, several interesting open-loop schemes have been proposed. These schemes assume that every client has a substantial amount of available storage capacity, which can be used to temporarily store portions of the movie. With such schemes, the server's transmission schedule and the algorithm used by the client to decide whether or not to "record" any given transmitted block jointly ensure that every segment of the movie will be on the user's disk by its viewing time. We next describe several such schemes. For facility of exposition and because that is the only case addressed by some of the schemes, we assume a fixed video rate. We conveniently use "length" to mean both viewing time and amount of data, and the video rate as our unit of data rate.

De Bey [3] suggested to partition a movie into $N=L/D$ segments of (fixed) length *D*, where *D* is the permissible viewing-commencement delay. The server transmits segment m, $1\leq m\leq N$, at the nominal video rate once every m time slots; the duration of a time slot equals *D*. (Multiple segments are transmitted concurrently during some of the time slots.) The client policy is to receive and store all the transmitted segments that have yet to be viewed. The server and client policies jointly ensure that every segment is available to the client by the time at which it is viewed. The **mean** transmission rate with this scheme for a movie of length *L* and permissible viewing delay *D* is

$$\overline{R_t}=\sum_{m=1}^{N}\frac{1}{m}\approx \ln\left(\frac{L}{D}\right).$$

This is a dramatic improvement over the Baseline scheme. However, the peak transmission rate is much higher than its average. Furthermore, this method imposes heavy requirements on the client hardware: recording rate reaches more than 12 times the video rate for $L/D=100$, and a client must be able to store nearly one half of the movie. With the Baseline scheme, recording rate equals the video rate and no storage is required. Finally, the direct coupling among viewing-commencement delay, movie length, and transmission policy prevents a trade-off between the consumption of various resources.

The *Pyramid* scheme [4] entails partitioning the movie into segments such that each segment is α times longer than the preceding one. All segments are then transmitted concurrently at the same data rate. (This achieves a similar, albeit not identical, effect as the assignment of different transmission rates to fixed-size segments.) At any given time, the client records data from at most two (consecutive) segments. The value of α is chosen according to the permissible delay and available transmission bandwidth, and to ensure that the client will always have the upcoming data available in its storage so as to prevent viewing glitches.

The *Permutation Pyramid* scheme [5] is based on the Pyramid. It creates multiple copies of each segment, partitions each copy into fixed size blocks, and interleaves the transmission of the copies such that same-numbered blocks are equispaced in the interleaved stream. The resulting stream is then transmitted at a fixed per-channel data rate. The result is a shorter viewing commencement delay as well as less client storage and recording bandwidth for the same aggregate transmission rate.

The *Harmonic broadcasting* scheme [6] uses segments of exponentially increasing size. The mean per-segment transmission rate is similar to De Bey's scheme. The advantage is a fixed (in time) aggregate transmission rate.

The *Staircase* scheme [7] presents yet another variant of the same general idea. Here, β channels, each with a transmission rate equal to the video rate, are used. The ith channel, $1\leq i\leq\beta$, is partitioned into 2^i subchannels. Next, the movie is divided into $N=2^\beta-1$ equisized segments (this is equal to the total number of subchannels). The segments are then assigned to the channels, one segment per subchannel, in ascending order of segments and channels. Finally, the segments assigned to each channel are interleaved with fine granularity such that the starting points of consecutive segments are staggered by equal distances, and the result is transmitted at the fixed channel rate. Transmissions occur concurrently on all channels. The client monitors the channels and records information that will not arrive again prior to its viewing time. This scheme outperforms the previous ones, except for the transmission bandwidth in the absence of client-storage and recording-rate constraints. In the latter case, the Harmonic scheme requires a lower transmission rate.

A close look at the NVOD problem reveals a multi-dimensional design space. All previously proposed schemes are constrained to highly regular designs, resulting in very limited flexibility. Moreover, there is no direct relationship between the design parameters in these schemes and the dimensions of the design space, so it is difficult to "navigate" intelligently even within the artificially constrained space.

Like the previous open-loop schemes, we also exploit client storage. Unlike the prior art in this framework, however, we employ a novel algorithmic approach to optimize the transmission scheme. In so doing, we can minimize transmission rate while directly satisfying constraints on the remaining dimensions. Alternatively, we can minimize a different parameter subject to constraints on the remaining ones and a given transmission rate.

The remainder of the paper is organized as follows. In Section 2, we analyze the data path and present the NVOD design space. In Section 3, we present some bounds and sketch our approach. Section 4 offers a comparison with prior art, and Section 5 offers concluding remarks.

2. The open-loop NVOD design space

In this section, we explore the NVOD design space, and derive two important design points: 1) no client resource requirements, and 2) minimum server transmission rate.

The provision of open-loop NVOD service to storage-capable clients entails (preferably) fixed-rate transmission by the server over a distribution network. Every client then selects the portions of the transmitted material that are relevant to it at any given time. A client "records" the selected material and plays it to the viewer at the right time. An NVOD solution thus comprises a transmission scheme, executed by the server, and a corresponding selection algorithm that is executed by each client.

The server's task entails reading data from disk and transmitting it over the distribution network. The required transmission rate for a single NVOD movie poses no problem to a high-performance disk drive and any PC. Moreover, in view of the large ratio of clients to NVOD movies, the cost of a server is not critical. Therefore, we do not consider server resources in the designs.

The bandwidth of the distribution network is clearly a critical resource, so aggregate transmission bandwidth R_t is an important dimension of the design space.

The demands placed on the user-premise equipment ("client") are of utmost importance. If it is dedicated to the NVOD service, cost reduction is critical due to the large quantities. If it is embodied in existing equipment, e.g., a computer, the NVOD requirements will determine the fraction of viewers to which the service can be offered. If temporary storage is provided by a disk drive, then disk bandwidth and required storage capacity are both critical client resources. If the storage medium is semiconductor memory, only capacity matters. So, the peak client recording rate, R_r, and its peak storage requirement, S_{max}, are two more dimensions. Once actual viewing begins, the client records data to disk while reading back data for playback. Thus, the maximum permissible recording rate is equal to the disk's effective rate minus the video rate.

Even if a disk drive is used as the client's primary medium for temporary storage, it is possible to use the client's memory to mitigate the recording-bandwidth requirements. One could therefore include the available amount of client memory as another design parameter. This use of memory, however, is closely related to a variety of implementation issues and is not addressed in any depth by the previously proposed schemes. In order to facilitate comparison and to focus on the main contributions of this paper, we do not discuss this use of memory in any detail.

Having established the dimensions of the design space, let us next derive two interesting optimal design points. One such point, optimized for *minimum consumption of* client resources, is provided by the Baseline scheme: $R_t=L/D$; $R_r=1$ (or zero, depending on definitions); $S_{max}=0$. We now derive the point that minimizes aggregate transmission rate R_t.

Consider a movie of length L seconds. We use m to refer to a location within the movie. Similarly, we use v to denote the elapsed time since the viewer requested to view a movie (and thus began to receive and record it). $v(m)$ denotes the time at which the viewer views location m in the movie. (The reader may think of dividing the movie into single-second segments, with $m=0,1,..L-1$ denoting the segment number.) The video rate is used as the unit of data rate.

Lemma 1: In any feasible open-loop NVOD transmission scheme, $r_t(m)>(1/v(m))$.

Proof: By contradiction. The amount of segment-m data that is transmitted during $v(m)$ seconds is $r_t(m)\cdot v(m)$. Thus, if $r_t(m)<(1/v(m))$, only part of the segment will be recorded in time for its viewing. □

Proposition 2: For any open-loop storage-assisted VOD system, the minimum aggregate transmission rate R_t

(expressed in units of the movie's data rate) required to satisfy a viewing-commencement delay *D* is

$$R_t^{\min} = \ln\left(1+\frac{L}{D}\right).$$

Proof: Follows from Lemma 1 by choosing the minimum values for $r_t(m)$, substituting $v(m)=D+m$, and approximating the summation over *m* with integration. □

Since we are allocating the minimum transmission rates, segment *m* must be recorded during the time interval $(0,v(m))$. The client's recording rate is therefore

$$r_r(v)=\begin{cases}\int\limits_{m=0}^{L}\frac{1}{m+D}dm=\ln\frac{L+D}{D}, & 0<v\le D\\ \int\limits_{m=v-D}^{L}\frac{1}{m+D}dm=\ln\frac{L+D}{v}, & D<v\le L+D\end{cases}$$

The required storage space equals the recorded data minus segments that have already been watched. For $v>D$,

$$S(v)=\int\limits_0^v r_r(t)dt-1\cdot(v-D)=v\cdot\ln\left(\frac{L+D}{v}\right).$$

Maximization of *S* over the viewing time yields

$$v(S_{\max})=\frac{(L+D)}{e}=0.37(L+D);$$

$$S_{\max}=0.37(L+D).$$

Also, the peak client recording rate is

$$R_r=\ln\left(\frac{L+D}{D}\right).$$

Remark. The result for S_{max} holds for $D<1.6L$, So $S_{max}<L$. Also, minimizing transmission rate requires that the client be able to store nearly 40% of the movie.

Since there is a feasible design with no client storage or recording requirements, it follows that any constraints on those can be satisfied. However, if their values are smaller than the ones obtained with the minimum transmission rate, the latter will have to be increased.

3. The Tailor-Made NVOD approach

Server action. The Tailor-Made scheme assigns a transmission rate $r_t(m)$ to the *m*th segment, and all segments are transmitted concurrently and repetitively by the server without any synchronization among their starting times. (In practice, blocks belonging to the various segments can be time-interleaved on one or several channels.) Thus, the server transmits at a fixed data rate $R_t=\sum r_t(m)$, regardless of whether the video rate is fixed or variable. The bandwidth allocation algorithms produce the values of $r_t(m)$ that minimize R_t while satisfying various client-related requirements and constraints.

Client action. The client's task, as before, is to record received information if and only if it is not scheduled to arrive again prior to the earliest time at which it may be needed for viewing. Recording of data packets constituting the contents of a given segment need not begin with the one containing the beginning of that segment.

If a movie segment is allocated a higher transmission rate than the minimum, the client may postpone the time at which it begins to record this segment. Since the storage space for any movie segment is released only after it has been viewed regardless of its recording time, it follows that minimization of client storage requirements favors the recording of any given segment at the latest possible time. With the unconstrained client storage consumption or its recording rate) minimum-transmission-rate solution, a client's recording rate decreases monotonically with viewing time, as segments that have been viewed are no longer being recorded. Therefore, postponing the commencement of recording certain segments can reduce the peak recording rate. This also suggests that it is best for the client to start recording any given segment as late as possible, and to record it continuously until its viewing time. Denoting the (viewing) time at which a client begins to record movie segment *m* by $s(m)$, it follows that

$$r_t(m)=\frac{1}{v(m)-s(m)}.$$

$$r_r(v)=\sum_{\substack{m:s(m)<v\\ v(m)>v}} r_t(m).$$

$$S(v)=\sum_{\substack{m:s(m)<v\\ v(m)>v}} r_t(m)\cdot\left(v-s(m)\right)$$

Based on the above observations, we have been able to algorithmically optimize the transmission-rate assignments to the various movie segments so as to minimize aggregate transmission rate while meeting constraints on peak client storage and/or peak client recording rate, and given the movie parameters and permissible viewing delay.

Support for pausing. Unlike with the Baseline scheme, pausing must be "active": the client must continue to record new data while deleting data that is guaranteed to be received again prior to the earliest time at which it may have to be presented. It can be shown that support of pausing does not increase resource consumption.

4. Results and comparison with prior art

The purpose of this section is to compare the new Tailor-Made designs with previously proposed open-loop schemes, and to provide several representative designs.

The comparison with other schemes is complicated by the fact that their results usually do not explicitly refer to all the dimensions of the design space (aggregate transmission rate R_t, viewing-commencement latency D, peak client recording rate R_r, and peak client storage consumption S_{max}). Also, they cannot tailor designs to exact specifications. Therefore, we were forced to reconstruct several design points for each scheme and to then tailor our design to match these points in all but one dimension. The plots presented in papers describing the previously-proposed schemes depict the value of one parameter, e.g., R_t, versus that of another such as D. These plots are somewhat misleading, however, because the remaining parameters are not held constant. All this makes the presentation of comparative plots either meaningless or prohibitively effort-consuming. Instead, we resort to tables with a small number of representative results. The Pyramid scheme [4] is dominated by the Permutation Pyramid [5]. The Harmonic scheme [6] performs well in terms of R_t in the unconstrained case, but even then it is inferior (even in terms of R_t) to the Tailor-Made designs. This is due to the coupling between segment length and viewing-commencement delay D. At its only design point for any given combination of L and D, the Harmonic scheme requires a large amount of client storage, rendering it impractical in many situations.

In view of the above, Tailor-Made will only be compared with the Permutation Pyramid and Staircase schemes. In order to conform to the format of the results for the other schemes, we cite the total disk rate R_d rather than its recording rate R_r. For Tailor-Made, we use $R_d \equiv R_r+1$, reflecting the fact that we did not exploit the extra bandwidth that is available for recording while $v<D$.

Tables 1 and 2 present transmission-rate comparisons between the Tailor-Made scheme and the Permutation Pyramid and Staircase schemes, respectively, for equal values of D, R_r and S_{max}. The advantage of the Tailor-Made designs is clearly evident. Another, equally important advantage, is the ability to tailor the design to specifications.

	D[sec]	S_{max} [%]	$R_d[R_v]$	$R_t[R_v]$
Perm. Pyr.	60	23.3	3.65	18.6
Tailor-Made	60	23.3	3.65	5.29
Perm. Pyr.	30	24.3	3.89	20.2
Tailor-Made	30	24.3	3.89	5.93
Perm. Pyr.	10	24.5	4.25	22.7
Tailor-Made	10	24.5	4.25	7.05

Table 1. Transmission-rate: Tailor-Made Vs. Permutation Pyramid. *L=120 min.*

	D[sec]	S_{max} [%]	$R_d[R_v]$	$R_t[R_v]$
Staircase	60	24.8	2.63	7
Tailor-Made	60	24.8	2.63	6.14
Staircase	30	24.9	2.75	8
Tailor-Made	30	24.9	2.75	6.83
Staircase	10	25.0	2.83	10
Tailor-Made	10	25.0	2.83	8.18

Table 2. Transmission-rate comparison: Staircase Vs. Tailor-Made. *L=120 min.*

In Tables 3 and 4, we focus on a single design point of the Permutation Pyramid and Staircase schemes, respectively. We use our approach to tailor a design that is constrained in all but one dimension to the same values as that of the referenced scheme, and compare the values in the remaining dimension. It can readily be seen that the Tailor-Made approach dominates the referenced schemes. (In the third row of Table 3, $D<1$sec was deemed impractical, hence the lower value of R_t.)

	D[sec]	S_{max} [%]	$R_d[R_v]$	$R_t[R_v]$
Perm. Pyr.	30	24.3	3.89	20.2
Tailor-Made	30	24.3	3.89	5.9
Tailor-Made	1	24.3	3.89	10.2
Tailor-Made	30	2.7	3.8	20.2
Tailor-Made	30	24.3	2.09	20.2

Table 3. Tailor-Made Vs. Permutation Pyramid. *L=120 min.*

	D [sec]	S_{max} [%]	$R_d[R_v]$	$R_t[R_v]$
Staircase	30	24.9	2.75	8
Tailor-Made	30	24.9	2.75	6.8
Tailor-Made	14	24.9	2.75	7.9
Tailor-Made	30	12.5	2.75	7.9
Tailor-Made	30	24.9	2.48	7.9

Table 4. Tailor-Made Vs. Staircase. *L=120 m.*

Next, we demonstrate the capabilities and performance of the Tailor-Made designs by showing several possible configurations for two video rates: 1.2Mb/s (MPEG-1 streams) in Table 5, and 4.8Mb/s (PAL broadcast quality MPEG-2 streams) in Table 6. In choosing the configurations, we targeted two communication fabrics: 10Mb/s Ethernet and a 25-30Mb/s cable channel.

L [min]	D[sec]	S_{max}[MB]	R_d[KB/s]	R_t [Mb/s]
30	10	32	766	8.43
30	30	32	596	7.13
120	10	265	412	10
120	30	162	412	9
120	30	64	710	13
120	30	32	596	22.35
120	30	12	650	27

Table 5. Tailor-Made design points for R_v=1.2Mbit/s (MPEG-1).

L [min]	D[sec]	S_{max} [MB]	R_d [KB/s]	R_t [Mb/s]
120	30	1,600	3,900	26.4
120	60	1,600	3,480	23.1
120	60	1,000	2,190	25.4

Table 6. Tailor-Made design points. R_v= 4.8Mbit/s (PAL broadcast-quality MPEG-2).

Referring to the bottom row of Table 5, for example, we see that even with only 12MB of client storage space, one can use a single 6MHz television channel (used with a cable modem) to offer a 120-minute MPEG-1 movie in NVOD mode with $D=30s$. This does not require a disk drive! Referring to the 5th row, two such movies can be offered concurrently if the client has 64MB of available memory, which will be very reasonable in the near future. Referring to the 4th row, 162MB of client storage would permit the provision of three such movies on a cable channel or one over 10Mb/s Ethernet. Tiny magnetic disk drives weighing only 20gr with double this capacity are becoming available. With 1.6GB of client storage, a single MPEG-2 movie with quality similar to broadcast-quality PAL can be offered over a cable channel. Recently, "digital TV" sets with multi-GB disk drives have been announced, so this configuration is also likely to become viable in the near future. Of course, the use of a PC as the platform for the client is likely to provide the required resources free of charge. We note in conclusion that the foregoing discussion also illustrated the importance of the ability to design the NVOD system to specifications.

Remark. The results in every row of the tables can be scaled in order to obtain numerous design points. To do so, R_t, R_r, R_v, and S_{max} must be changed by the same multiplicative factor while keeping D and L unchanged.

Finally, it should be noted that the results presented in this section for the Tailor-Made scheme are conservative. No use was made of RAM buffers to mitigate the disk-rate requirement, and we even did not take advantage of the fact that as long as $v<D$, no data is read from disk so its entire bandwidth can be devoted to recording.

5. Conclusions

This paper showed that the NVOD schemes that restrict themselves to regular structures (e.g., segment sizes that form a geometric sequence) are both substantially sub-optimal and inflexible. The gap is closed by the Tailor-Made algorithmic approach, which can moreover be used with both fixed and variable video rates. It is important to observe that, regardless of the behavior of the video rate, the server always transmits at a fixed data rate. Also, this is an open-loop scheme, so it is truly scalable to an unlimited number of viewers. Finally, the numbers are very attractive for the new digital broadcast-oriented distribution networks such as cable and satellite.

A prototype that demonstrates the Tailor-Made designs has been built, using PCs as both server and clients. The algorithms for assigning transmission rates to segments have been coded. Their computational complexity is approximately quadratic in the number of segments into which the movie is divided. Various further reductions in complexity are possible, but are not critical because these algorithms are normally executed off-line.

Topics of ongoing research include utilization of RAM buffers to reduce the required disk bandwidth, efficient error-correction, and extension of the approach to concurrently provide different qualities of service to different classes of clients based on their capabilities. The approach is also being adapted to configurations in which a single client cannot "listen" to all the transmitted data. This applies to point-to-point networks such as ADSL as well as to the case of limited bandwidth in the client's communication equipment. Finally, an extension that integrates the transmission of a live event with a smooth transition to NVOD is being developed.

Acknowledgment. Hilit Grosberg, Liane Lewin, Yona Sivan and Yair Wiener took part in the implementation of the prototype. Michael David, Evgeny Rivkin and Tomer Kol provided assistance.

6. References

[1] A. Dan, P. Shahabuddin, D. Sitaran and D. Towsley, "Channel allocation under batching and VCR control in Movie-On-Demand servers", Technical report, IBM Research, 1994.

[2] L. Golubchik, L. C.-S. Lui, and R. R. Muntz, "Reducing I/O Demand in Video-On-Demand Storage Servers", Proceedings of ACM SIGMETRICS, pp. 25-36, 1995.

[3] H. C. De Bey, "Program transmission optimisation", United States Patent Number 5,421,031, Mar 1995.

[4] S. Viswanathan and T. Imielinski, "Metropolitan area video-on-demand service using pyramid broadcasting", Multimedia Systems, vol. 4, pp. 197-208, 1996.

[5] C.C. Aggarwal, J.L. Wolf, and P.S. Yu, "A permutation-based pyramid broadcasting scheme for video-on-demand systems", IEEE Proceedings of the International Conference on Multimedia Computing and Systems, pp. 118-126, Jun 1996.

[6] L-S. Juhn and L-M. Tseng, "Harmonic broadcasting for video-on-demand service", IEEE Trans. Broadcasting, vol. 43, no. 3, pp. 268-271, Sep 1997.

[7] L-S. Juhn and L-M. Tseng, "Staircase data broadcasting and receiving scheme for hot video service", IEEE Trans. Consumer Electronics, vol. 43, no. 4, pp. 1110-1117, Nov 1997.