Probabilistic Assignment of Movies to Storage Devices in a Video-On-Demand System[1]

T.D.C. Little and D. Venkatesh

Multimedia Communications Laboratory
Department of Electrical, Computer and Systems Engineering
Boston University, Boston, Massachusetts 02215, USA
(617) 353-9877, (617) 353-6440 fax
{*tdcl,dinesha*}*@bu.edu*

MCL Technical Report 11-01-1993

Abstract–A video-on-demand server must satisfy a large customer base and a diverse archive of movies under changing movie popularity and daily load peaks. These requirements must be satisfied under the constraints imposed by storage device costs, capacities, I/O bandwidths, and geographic locations.

In this paper we describe a partitioning of video data (movies) onto a video-on-demand storage hierarchy to achieve efficient storage and I/O bandwidth utilization. Our approach uses a probabilistic model of movie popularity in data distribution and replication to balance user requests with available disk I/O bandwidth. The results can be applied in the design of a distributed video-on-demand system.

Keywords: Multimedia, video-on-demand, video servers.

---

[1] In *Lecture Notes in Computer Science*, Vol. 846, (Proc. 4th International Workshop on Network and Operating System Support for Digital Audio and Video), D. Shepherd, G. Blair, G. Coulson, N. Davies and F. Garcia, Eds., Springer-Verlag, October, 1994. This work is supported in part by the National Science Foundation under Grant No. IRI-9211165.

1 Introduction

Video-on-demand (VOD) describes a new form of information delivery that is causing a great commotion in the industries of cable TV, telecommunications, personal computers, and software. It is a multimedia application with tremendous potential commercial value.

Basic features of the simplest VOD systems are services to allow selection, delivery, and viewing of movies through interaction with a presentation device such as a television and remote control. VOD also supports applications such as interactive learning, catalog shopping, and general information browsing. In its basic form, a VOD interaction scenario consists of a movie database perusal or query, a request for a feature presentation, and movie viewing. VOD interaction capabilities can be classified into four types [5]. These are broadcast (No-VOD), pay-per-view (PPV), quasi video-on-demand (Q-VOD), and true video-on-demand (T-VOD), representing the full spectrum of interaction options. PPV services are similar to those available on existing cable-TV systems. The user has little control over the playback, and must receive the feature presentation at predetermined times set by the service provider. Q-VOD services allow the user limited control over the viewing by offering staggered movie start times. The additional sessions allow viewers to jump from session to session to gain access to a different portion of the presentation. T-VOD dedicates an entire session[2] to a single user and provides individualized user control over the presentation. T-VOD is the most difficult service to provide.

The main issues in designing a VOD system are primarily related to resource allocation. Issues such as disk partitioning and bandwidth reservation are crucial for supporting a VOD system [2, 3, 6]. A typical architecture for a VOD system consists of a set of video databases (VDBs) interconnected to a set of servers which perform routing of VOD traffic to individual users via an interconnection network. Users connect to the VOD service through network access points and customer premises equipment (CPE). In this model, it is envisioned that each server manages connection-oriented VOD traffic for multiple sessions.

Various studies have addressed specific issues in the design of a VOD system. The general architectural considerations in designing a VOD system have been investigated by Gelman et al. [5], Loeb [9], Sincoskie [14], and Sutherland and Litteral [15]. Chen and Little [2], Rangan et al. [12], and Christodoulakis and Faloutsos [3] have investigated physical disk organizations necessary for supporting VOD systems. Little et al. [8] describe a prototype

---

[2] A *session* represents the I/O bandwidth necessary to support the delivery of a single movie. *Users* and *sessions* are distinct because multiple users can participate in a single session (e.g., PPV or Q-VOD).

VOD system supporting content-based queries on on a video database for VOD. Ramarao and Ramamoorthy [11] use a probabilistic model for the assignment of videos to storage assuming a three tiered memory hierarchy. Practical system considerations are addressed by Anderson and Homsy [1], Lougher and Shepherd [10], and Silvers and Singh [13]. Commercial applications of VOD are currently being pursued by numerous organizations.

Performance considerations for VOD system design include the impact of user behavior, video-application bandwidths, storage subsystem I/O bandwidths, and video popularities. Although many of these issues have been studied independently, they have not been applied to the problem of efficient video data partitioning in a distributed VOD scenario.

In this paper, we propose a scheme that allows efficient utilization of disk storage space and I/O bandwidth in the context of an overall VOD system. We use a popularity model developed by Ramarao and Ramamoorthy [11]; however, we consider user behavior, disk I/O bandwidth constraints, and the distribution of movies to multiple storage devices.

The remainder of this paper is organized as follows. In Section 2, we introduce basic considerations for movie distribution. In Section 3, we describe our storage approach. Simulation results are presented in Section 4. Section 5 concludes the paper.

2 VOD System Characteristics

A VOD system must deliver interactive video services to a large customer base at a competitive cost relative to broadcast services. However, the complexity of a T-VOD session is much greater than a broadcast session due to the requirement for a mapping of individual users to sessions. In this section we examine the economics of a T-VOD system and introduce basic symbols used in our movie[3] mapping solution.

2.1 VOD System Economics

The basic disk storage economics for a VOD system are shown in Table 1. Here we consider the installation costs for a VOD system supporting 1000 sessions, a capacity for 1000 movies of 100 minute duration, and a connection-oriented I/O data rate of 1.5 Mbits/s. It is apparent that a larger number of storage devices yields a larger number of independent sessions as

---

[3] For the remainder of the paper we describe any long-lived, stored audio/video data as a "movie" without loss of generality.

Table 1: Disk storage economics

Size	**Cost/MB**	**Disks**	**Sessions/Disk**	**Cost/Session**	**Total Cost**
5 Gbytes	$ 1	225	4.4	$ 1,225	$ 1,125,000
25 Gbytes	$ 0.20	45	22	$ 225	$ 225,000

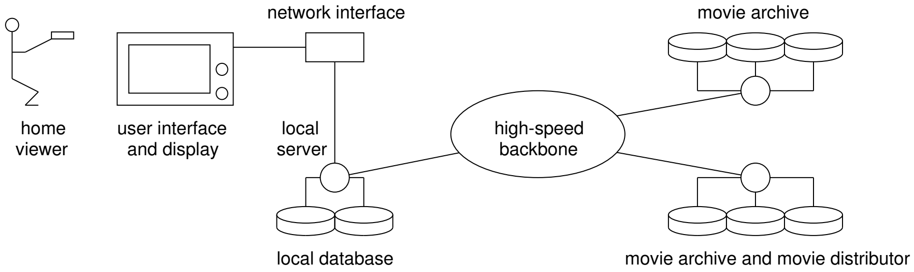

Figure 1: VOD system architecture

measured in available bandwidth. A larger number of T-VOD sessions and improved user interaction capabilities are provided by achieving aggregate storage capacity with many smaller disks. Table 1 illustrates the rough estimates on the cost for two storage device capacities. However, this basic analysis assumes a perfect mapping of user movie choices to available disk drives, i.e., a selected movie can be assigned to the free I/O bandwidth of any disk. In reality, user movie choices must be mapped to sessions from devices which store the selected movies. If an imperfect mapping is used, the limitations in disk I/O bandwidth prevent this perfect disk utilization. In the following section we consider an architecture suitable for a reconfigurable mapping of movies to storage devices.

2.2 VOD System Architecture

We use a VOD architectural model as illustrated in Fig 1. Our proposed movie-disk distribution scheme is outlined as follows. On a daily basis a centralized movie distributor calculates the current movie popularities. From these popularities the movies are replicated as necessary (Section 3) and distributed during off-peak hours to localized storage systems. The movies are then available for viewing during peak hours based on anticipated demand.

Although centralized control is assumed for movie redistribution, the archival of movie "originals" can be performed in either a distributed or centralized fashion. This archive can

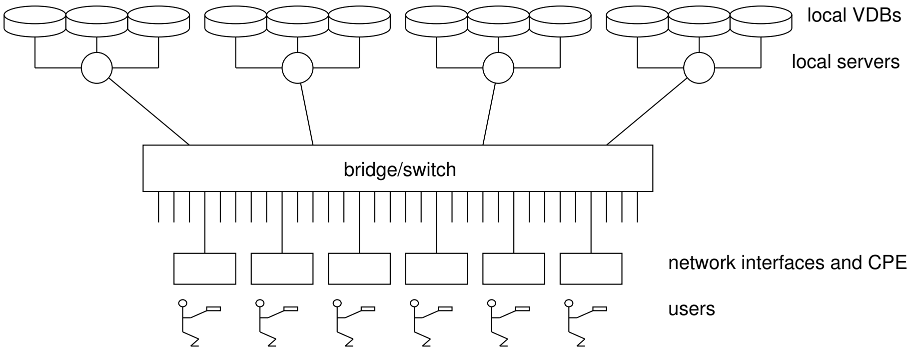

Figure 2: Local database architecture for the VOD system

be viewed as a very large database containing all available movies and stored on media with varying access latencies and I/O bandwidths.

This model is not unlike the multi-tiered memory architecture of computer systems. In our current scheme, a user's movie selection is limited to the set of movies available at the local database after movie distribution. We therefore concentrate on the data distribution problem and the design of the memory hierarchy at the local VDB.

The number of customers using consumer No-VOD (boadcast) systems varies over a 24 hour period. This "load" is maximum during the evening peak hours and slacks off during the night. In our T-VOD system the slack times will be used by the system to update the database organization and transfer movies back (and forth) from the central database.

In a typical VOD scenario, N customers access a local database containing K movies. The K movies are replicated to distribute access demand to sufficient I/O bandwidth depending on their popularity, and are stored on d disks. Such a system can be visualized as shown in Fig. 2. We restrict the number of movies K to be constant over a period of 24 hours (i.e., between movie redistribution times).

A reasonable partitioning approach for movies to a memory organization is based on movie popularity [11]. For example, movie popularity can be defined by a vector

$$P = [p_1, p_2, ..., p_K]$$

where

$$\sum_{i=1}^{K} p_i = 1.$$

It is not clear to us the exact form of the popularity distribution and we presently assume a form amenable to analysis. Furthermore, we assume this distribution to be stationary, i.e., the popularity distribution on the set of all known movies changes little over time. In contrast, the mapping of movies to popularities can change rapidly, but we assume that this mapping to popularities remains constant within an interval of 24 hours. This assumption is consistent with a recomputation of movie popularity on a daily basis.

For example, consider a VDB with $K = 10$. Let $[p_1, p_2, ..., p_{10}]$ be the popularities of the individual movies in descending order of popularity. A new movie which is most popular would map to p_1. As time passes, its popularity drops, until it needs to be discarded, being replaced by other movies. We assume here, that the vector $[p_1, p_2, ..., p_{10}]$ remains fairly constant over time, and this is the basis for our time-invariant distribution assumption.

We envision the VDB as having a finite capacity for movies. The VDB is an array of d disks, disk j holding k_j movies. Furthermore, each disk can only support a finite number of sessions, l, and therefore limits the peak number of active T-VOD users to l.

When device I/O bandwidth constraints are considered (i.e., limited sessions per device), we satisfy movie popularity demands by movie replication across the set of disks comprising the VDB. Therefore, the the total number of movies K is constrained by[4]

$$K \leq \sum_{j=1}^{d} k_j$$

We then define a replication factor $r_{ij} \in \{0, 1\}$, with a non-zero value indicating the replication of movie i on disk j.

The probability of accessing each movie copy is then computed as

$$p_{ic} = \frac{p_i}{N_{ci}} \tag{1}$$

where N_{ci} are the number of copies of movie i. Equation 1 implies that each copy of the movie is equiprobably accessible. As each disk can support only a finite number of sessions l, there is a lower limit on N_{ci}, given by

$$N_{ci} \geq \frac{p_i \times N}{l}$$

---

[4] Here we assume that no movie replication is allowed on the same device.

Table 2: Summary of symbols used in movie-disk assignment

K	Number of stored movies
N	User population
d	Number of storage devices
$N_i(t)$	Number of users accessing movie i at time t
N_{ti}	Number of copies of movie i
p_i	Popularity of movie i (probability of movie access)
P	Popularity vector $[p_1, p_2, ..., p_K]$
p_{max}	Popularity of the most popular movie
R_i	Replication vector for a movie i $[r_{i_1}, r_{i_2}, ...r_{i_d}]$
s_i	Probability of access for disk i
S	Probability vector for disk access $[s_1, s_2, ..., s_d]$
l	Number of sessions supported by a disk (homogeneous)
Q	Probability of available sessions from a disk
Z	$K \times d$ Population replication matrix
λ	Request arrival rate at the VOD system
P_B	Probability that a users call is blocked
P_{Bi}	Probability that a call is blocked at disk i
μ	Average call holding time

With a number of disks corresponding to N_{ci}, the session requirements are satisfied when all disks carrying this copy of the video are serving only this single movie. These parameters and their descriptions are summarized in Table 2.

3 Movie–Disk Assignment

We now introduce our proposed movie-disk assignment scheme. First we consider the computation of the movie popularity and then a proposed scheme for allocating movies to disks when there is an I/O bandwidth constraint limiting the number of sessions per device.

We assume that a daily computation of the popularity distribution is sufficient to track its dynamics. To compute this distribution, we use a linear predictor which takes into account the popularity of the movie over the previous two days. Let $N_i(t-2)$ and $N_i(t-1)$ be the number of users using the movie i on days $(t-2)$ and $(t-1)$ respectively. Then, using a linear predictor, the number of users expected to request movie i for day t is,

$$N_i(t) = 2N_i(t-1) - N_i(t-2), \quad N_i(t) \geq 0 \tag{2}$$

In other words, the number of users requesting the movie is a linear function of the number of users accessing it on the previous day. The slope is determined by the change in the number of users over the previous two days. This simple predictor can easily be replaced by a more complicated nonlinear predictor that considers more complex models of user behavior.

The popularity of a given movie i is then computed as

$$p_i = \frac{N_i(t)}{\sum_{m=1}^{K} N_m(t)} \tag{3}$$

When a new movie becomes available for the first time, it is given a predetermined value N_{new} based on prior estimates of the type of movie (e.g., home movie release, sporting event, music video). If there are N users, the number of sessions required for supporting the playout of a given movie i at any time is $N \times p_i$. If a disk can support l sessions, the minimum number of disks required to support the ith movie is given by

$$d \leq \left\lceil \frac{N \times p_i}{l} \right\rceil$$

In the limiting case, this condition implies placing a single copy of a movie onto a single disk. The disk array then supports the requisite number of sessions. If additional (different) movies are placed on the same disks then the probability of a successful connection decreases for a constant total disk space. This probability is minimal when all movies are placed on a single disk. Thus there is a tradeoff between the storage capacity of large disks and the I/O capacity of multiple smaller disks. In addition, we can show that there is little gain when multiple copies of a movie are placed on the same disk as indicated by the following conjecture.

Conjecture 1 *There is no improvement in connection setup probability when more than one copy of the same movie is placed on a video disk.*

Justification: Since we are assuming T-VOD, time separation in movie start times yields independent disk head movement as would occur for replicated movie copies on the same device. Hence, there is no gain by placing more than one copy of the same movie on the same disk. There could be marginal improvement due to reduced latency, but this would not be predictable due to the large variation in accesses from user to user. A priori knowledge of user interaction would be necessary during physical disk layout in order to achieve these gains.

Conjecture 1 places a lower bound on the number of disks required in the VOD system. If p_{max} is the probability of the most popular movie, the lower bound on the number of disks required in the system is,

$$d \geq \left\lceil \frac{N \times p_{max}}{l} \right\rceil$$

Violation of this bound by movie replication on the same disk leads to storage inefficiencies without I/O bandwidth gain. We therefore define the bounds for the number of disks required for the system, d as

$$\left\lceil \frac{N \times p_{max}}{l} \right\rceil \leq d \leq \sum_{i=1}^{K} \left\lceil \frac{N \times p_i}{l} \right\rceil$$

or

$$\left\lceil \frac{N \times p_{max}}{l} \right\rceil \leq d \leq \left\lceil \frac{N}{l} \right\rceil \tag{4}$$

The lower bound is obtained from the constraints imposed by the most popular movie, and the upper bound is derived from the extreme case when each disk holds only a single movie. These bounds are illustrated in Fig. 3 for $l = 5$, and various values of N. Fig. 3(a) demonstrates the bounds for a small number of sessions, while Fig. 3(b) illustrates the bounds for a large number of sessions.

We then define a popularity-replication matrix Z as a $K \times d$ element matrix such that

$$\sum_{i=1}^{K} z_{i,j} = s_j$$

$$\sum_{j=1}^{d} z_{i,j} = p_i$$

and

$$\sum_{i=1}^{K} \sum_{j=1}^{d} z_{i,j} = 1.$$

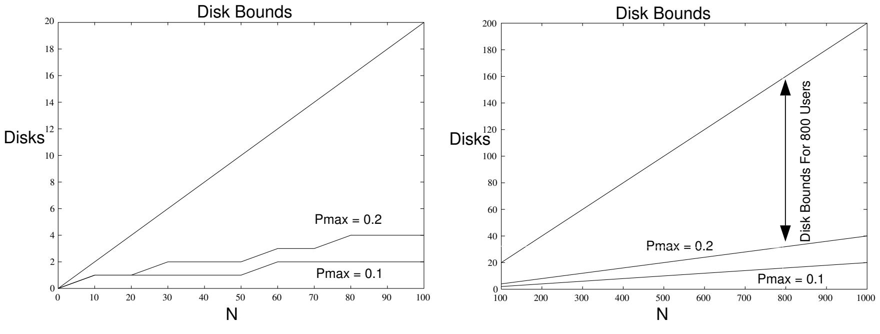

(a)

(b)

Figure 3: Movie-disk assignment bounds for VOD for (a) 0–100 users, and (b) 100–1000 users.

Here, s_j is the popularity of disk j being accessed by any user. Row i of Z corresponds to the product of R_i with p_{ic}. The probability of a successful connection Q from a disk is then defined as

$$Q = P[\text{ less than } l \text{ connections on the disk}].$$

To ensure that a maximum number of sessions are supported, Q must be as small as possible. Furthermore, the probability that a random request from a user results in a successful connection with the system must be maximized. We desire to allocate movies to disks such that the chance of a successful session is maximum, and we conjecture as follows.

Conjecture 2 *The probability that a user's request results in a successful connection is maximum when the movies are distributed such that each disk has a uniform probability of being accessed (i.e., $s_p \approx s_q$ for all p and q).*

Justification: The above conjecture is proved by showing that for any incoming request, the probability of blocking is a minimum when the movie-disk partition is uniform. The proof is as follows. Let requests for movies arrive into the system with a Poisson arrival rate λ. We are interested in the probability that this call is blocked, and want to minimize it. Given the popularity vector for the disk array $[s_1, s_2, ..., s_d]$, and assuming random requests, the arrival rate at any disk i, λ_i is given by

$$\lambda_i = s_i \lambda.$$

If we assume that the call holding times are exponentially distributed with a parameter μ, then the probability of blocking at the server i, P_{Bi} is given by Erlang's B Formula [7],

$$P_{Bi} = \frac{(\frac{s_i\lambda}{\mu})^l/l!}{\sum_{j=0}^{l}(\frac{s_i\lambda}{\mu})^l/j!} \tag{5}$$

The probability of blocking for a request to the system can be expressed as

$$P_B = \sum_{i=0}^{d} s_i P_{Bi}.$$

Substituting equation 5 for P_{Bi}, we obtain,

$$P_B = \sum_{i=0}^{d} s_i \frac{(\frac{s_i\lambda}{\mu})^l/l!}{\sum_{j=0}^{l}(\frac{s_i\lambda}{\mu})^l/j!}$$

or,

$$P_B = \frac{(\frac{s_i\lambda}{\mu})^l/l!}{\sum_{j=0}^{l}(\frac{s_i\lambda}{\mu})^l/j!} \sum_{i=0}^{d} s_i^{l+1} \tag{6}$$

The first part of the equation is a constant for a given λ. Therefore, P_B is a minimum when $\sum_{i=1}^{d} s_i^{l+1}$ is a minimum. It can readily be shown by the method of Lagrangian multipliers [4] that

$$min(P_B) \Rightarrow s_i = 1/d, \quad \forall i,$$

justifying our conjecture that the chance that a user's request for a session is successful is maximum when the probability of access of each disk is equal. However, this statement says little about the sensitivity of the system to a miss-ordering during movie-disk assignment. In the following section, we present results of simulations which address this sensitivity.

4 Simulation Results

In the previous section, we developed a scheme for allocation of movies to disks in a VOD system. In this section, we analyze the scheme with respect to sensitivities to sub-optimal movie-disk assignment and show simulation results.

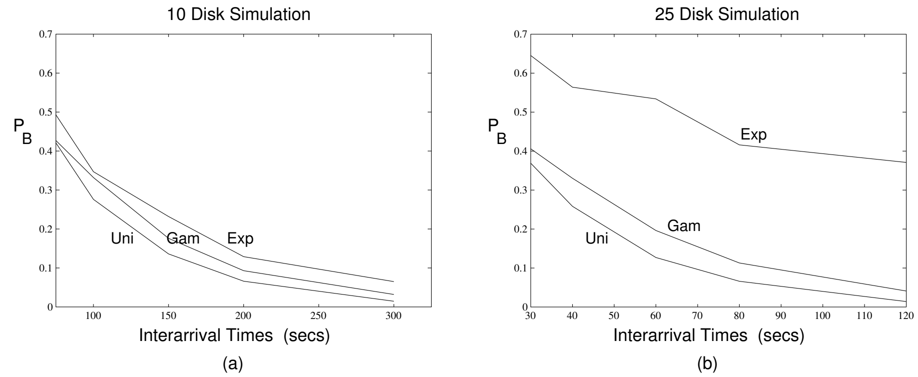

Figure 4: Blocking probabilities for movie popularity assignments for (a) 10 disks, and (b) 25 disks

The proposed scheme is an efficient strategy for the placement of data in a homogeneous system i.e., when all the disk servers are alike. To address the sensitivity of the scheme to a sub-optimal movie-disk assignment, we performed two simulations. The first simulated a VDB with 10 disks each containing 5 movies and supporting a maximum of 5 concurrent sessions. The second simulation was a variation of the first with 25 disks but was otherwise the same. Each session is assumed to have an exponentially distributed viewing time with a mean of 90 minutes. In addition, the two VDBs were studied under three different movie popularity assignments as follows.

1. A "uniform" ordering where all disks were equiprobably accessible (our proposed movie-disk assignment),
2. an "exponential" ordering, in which the movie popularities were exponentially distributed to disks, and
3. a "gamma" ordering in which the disk popularities were given a gamma distribution with parameter 2.

The exponential ordering results in some movie disk assignments which were excessively popular, while others were very unpopular. This effect is less acute for the gamma assignment policy. The results of these simulations are shown in Fig. 4 which shows the probability of new session blocking versus request interarrival times.

The results indicate that a uniform ordering of disks to popularities always yields the highest session availability. Furthermore, even a moderately "flatter" ordering such as the gamma, with a smaller variance as compared to the exponential, results in significant gains over the exponential ordering. These gains become more significant with increasing VDB sizes.

5 Conclusion

In this paper we have proposed an approach to replicating and distributing movies to disks in a video-on-demand system. Our scheme uses a probabilistic model based on movie popularity and considers the effect of limited disk I/O bandwidth on system architecture. Results indicate the bounds on the number of disks required in a VOD system for a given movie popularity vector and disk I/O bandwidth characteristic. The results also establish the conditions for a movie-disk assignment that yields the minimum probability of new session blocking.

The results of this study are best suited for rewritable storage systems with sufficient I/O bandwidth to support multiple video delivery sessions. Such devices can satisfy the local VDB requirements of daily rewrites and support for a large number of independent sessions. Devices with poor I/O bandwidth or read-only characteristics are more appropriate for movie archival.

There are a number of issues that we will investigate in the future. These include the implications of (1) customer pricing on the degree of interaction and session blocking; (2) the availability of network bandwidth for sessions and movie redistribution; (3) the transition of T-VOD sessions to Q-VOD sessions for service scaling; (4) consideration of heterogeneous disk capacities and I/O bandwidths; and (5) consideration of heterogeneous session characteristics.

References

[1] Anderson, D.P., Homsy, G.: A continuous media I/O server and its synchronization mechanism. Computer **24** (1991) 51-57

[2] Chen, H.-J., Little, T.D.C.: Physical storage organizations for time-dependent multimedia data. To appear at ACM FODO'93 (1993)

[3] Christodoulakis, S., Faloutsos, C.: Design and performance considerations for an optical disk-based, multimedia object server. Computer **19** (1986) 45-56

[4] Danø, S.: Nonlinear and dynamic programming. Springer Verlag Ltd. (1975)

[5] Gelman, A.D., Kobrinski, H., Smoot L.S., Weinstein, S.B.: A store-and-forward architecture for video-on-demand service. Proc. ICC (1991) 27.3.1-27.3.5

[6] Gelman, A.D., Halfin, S.: Analysis of resource sharing in information providing services. IEEE Global Telecommunications Conference and Exhibition Conference Record (1990) 312-316

[7] Kleinrock, L.: Queueing systems vol. 1. John Wiley and Sons (1975)

[8] Little, T.D.C., Ahanger, G., Folz, R.J., Gibbon, J.F., Reeve, F.W., Schelleng, D.H., and Venkatesh, D.: A digital video-on-demand service supporting content-based queries. ACM Multimedia'93 (1993) 427-436

[9] Loeb, S.: Delivering interactive multimedia documents over networks. IEEE Communications Magazine **30** (1992) 52-59

[10] Lougher, P., Shepherd, D.: The design and implementation of a continuous media storage server. Proc. 3rd Intl. Workshop on Network and Operating System Support for Digital Audio and Video (1992) 63-74

[11] Ramarao, R., Ramamoorthy, V.: Architectural design of on-demand video delivery systems: The spatio-temporal storage allocation problem. Proc. ICC'91 (1991) 17.6.1-17.6.5

[12] Rangan, P.V., Vin, H.M., Ramanathan, S.: Designing an on-demand multimedia service. IEEE Communications Magazine **30** (1992) 56-64

[13] Silvers, Singh, E.: Multimedia communications on the NYNEX shuttle. Proc. COMPCON (1992) 84-87

[14] Sincoskie, W.D. : System architecture for a large scale video on demand service. Computer Networks and ISDN Systems **22** (1991) 155-162

[15] Sutherland, J., Litteral, L. Residential video services. IEEE Communications Magazine **30** (1992) 36-41