

# Throughput-Competitive Admission Control for Continuous Media Databases

Minos N. Garofalakis\*

University of Wisconsin–Madison

*minos@cs.wisc.edu*

Yannis E. Ioannidis<sup>†‡</sup>

University of Wisconsin–Madison

*yannis@cs.wisc.edu*

Banu Özden

Bell Laboratories

*ozden@research.bell-labs.com*

Avi Silberschatz

Bell Laboratories

*avi@research.bell-labs.com*

## Abstract

Multimedia applications require a guaranteed level of service for accessing Continuous Media (CM) data, such as video and audio. To obtain such guarantees, the database server where the data is residing must employ an admission control scheme to limit the number of clients that can be served concurrently. We investigate the problem of *on-line admission control* where the decision on whether to accept or reject a request must be made without any knowledge about future requests. Employing *competitive analysis* techniques, we address the problem in its most general form with the following key contributions: (1) we prove a tight upper bound on the competitive ratio of the conventional Work-Conserving (WC) policy, showing that it is within a factor  $\frac{1+\Delta}{1-\rho}$  of the optimal clairvoyant strategy that knows the entire request sequence in advance, where  $\Delta$  is the ratio of the maximum to minimum request length (that is, time duration), and  $\rho$  is the maximum fraction of the server's bandwidth that a request can demand; (2) we prove a lower bound of  $\Omega(\frac{\log \Delta}{1-\rho})$  on the competitive ratio of *any deterministic or randomized* admission control scheme, demonstrating an exponential gap between greedy and optimal on-line solutions; (3) we propose simple deterministic schemes based on the idea of *bandwidth prepartitioning* that guarantee competitive ratios within a small constant factor of  $\log \Delta$  (i.e., they are *near-optimal*) for sufficiently large server bandwidth; (4) we introduce a novel admission control policy that partitions the server bandwidth based on the expected popularities of different request lengths and present a set of preliminary experimental results that demonstrate the benefits of our policy compared to WC. We believe that our results offer new insights to other optimization problems that arise in CM data management, including data placement and load balancing in distributed CM databases.

\*Part of this work was done while the first author was visiting Bell Laboratories.

<sup>†</sup>Author's present address: Department of Informatics, University of Athens, Hellas (Greece).

<sup>‡</sup>Partially supported by the National Science Foundation under Grant IRI-9157368 (PYY Award), and by grants from IBM, DEC, HP, AT&T, Informix, and Oracle.

## 1 Introduction

Next generation database systems will need to provide support for various forms of multimedia data such as images, video, and audio. These new data types differ from conventional alphanumeric data in their characteristics, and hence require different techniques for their organization and management. A fundamental issue is that digital video and audio *streams* consist of a sequence of media quanta (video frames or audio samples) which convey meaning only when presented continuously in time. Hence, a multimedia database server needs to provide a guaranteed level of service for accessing such *continuous media* (CM) streams in order to satisfy their pre-specified real-time delivery rates. Given the limited amount of computing resources (e.g., memory, disk bandwidth, disk storage), it is a challenging problem to design effective resource management algorithms that can provide on-demand support for a large number of concurrent continuous media clients.

Typically, clients issue requests for the playback of specific CM *clips* (i.e., contiguous portions of audio or video) residing at the database server. A crucial component of CM service is the *admission control* mechanism, which is invoked whenever a new request arrives to decide whether to accept or reject the request. By accepting a request, the server commits to satisfy the resource requirements of the corresponding playback stream throughout its execution, whereas rejected requests must pursue a different course of action (depending on the application)<sup>1</sup>. The effectiveness of the admission control component is of vital importance for the following reasons. First, the resource requirements of CM applications are high. Second, they require fractions of the server's resources to be reserved to meet their stringent performance requirements. Third, these applications tend to last for relatively long periods of time. Reserving large portions of the resources for long durations can result in drastic degradation of server utilization if the server makes wrong decisions whom to admit.

An important characteristic of admission control is the introduction of an *on-line decision making* element – the decision of whether to accept or reject a request has to be made without any knowledge of future requests, with the understanding that once a request is accepted, it is guaranteed a level of service throughout its duration (i.e., the schedule is *non-preemptive*). In this paper, we study the implications of the on-line nature of the problem which has,

<sup>1</sup>Our model corresponds to the *Full-VOD* service model [1, 8, 22]. Other service models have also been explored in the literature [17].

for the most part, been ignored in the multimedia literature. Our performance metric for admission control strategies is the *total server throughput* over a sequence of requests and our methodology is based on the competitive analysis framework for on-line algorithms [30]. The basic quality metric in this framework is the *competitive ratio* of an on-line algorithm, which is defined to be the maximum (over all possible request sequences) value of the ratio of the performance of the optimal off-line algorithm for a request sequence to the performance of the on-line algorithm for the same request sequence. Note that, by definition, competitive analysis is tantamount to a worst-case analysis in the off-line case. An algorithm with a low competitive ratio is one that performs close to optimal in all situations. Since no assumptions are made about the sequence of requests offered to the server, the competitive ratio provides a very *robust* measure of performance.

We assume a centralized database server where incoming playback requests require some fraction of the server's bandwidth for some period of time. For example, a request to view a half-hour MPEG-1 video clip requires 1.5 Megabits per second (Mbps) of the server's bandwidth for the 30 minutes of playback. We consider two different cases of the problem. In the first case, we assume that all requests require the same fraction of the server's bandwidth (e.g., all clips are MPEG-1 encoded videos); thus, the server can be viewed as a collection of available playback channels. In the second, more general case, different fractions of the server's bandwidth can be reserved. We show that the conventional Work-Conserving (*WC*) policy where an incoming request is always admitted if there is sufficient bandwidth to accommodate it, can behave poorly in an on-line setting. More specifically, we show that the competitive ratio of *WC* is  $1 + \Delta$  for the case of identical bandwidth requests and  $\frac{1+\Delta}{1-\rho}$  for the case of variable bandwidth requests, where  $\Delta$  is the ratio of maximum to minimum request length and  $\rho$  is the maximum fraction of the server's bandwidth that a request can demand. We introduce novel admission control strategies based on the idea of *prepartitioning* the bandwidth capacity of the server among requests of different length and prove that, for sufficiently large server bandwidth, these strategies are  $O(\log \Delta)$ -competitive. We also show an  $\Omega(\log \Delta)$  (resp.  $\Omega(\frac{\log \Delta}{1-\rho})$ ) lower bound on the competitive ratio of any deterministic or randomized algorithm for the identical (resp. variable) bandwidth case, thereby establishing the near-optimality of our on-line algorithms. Based on the above results, we propose a bandwidth prepartitioning scheme that makes use of clip popularities to ensure good average-case as well as worst-case performance. The results of our preliminary experimental study verify the benefits of our scheme as compared to *WC*. More specifically, both algorithms are shown to perform adequately well when the server is underutilized or persistently overloaded. However, we expect that a well designed system has undergone effective capacity planning and, therefore, will not be overloaded persistently but only at short time intervals. We capture such short term overloads in our experiments, and demonstrate that our admission control scheme outperforms *WC* substantially under these workloads.

The remainder of this paper is organized as follows. Section 2 reviews related work in the area of multimedia databases and on-line algorithms. In Section 3, we provide the necessary definitions and formulate the on-line admission control problem. Section 4 introduces our novel Bandwidth Prepartitioning strategies and presents our competitive analysis results for both identical bandwidth and vari-

able bandwidth requests. In Section 5 we discuss the findings of a preliminary experimental study with the admission control schemes described in this paper. Finally, Section 6 concludes the paper with a discussion on the implications of our results in the context of data placement in distributed CM databases. All theoretical results in this paper are stated without proof due to space constraints. The details can be found in the full version of this paper [16].

## 2 Related Work

Resource scheduling issues in CM databases have attracted considerable interest from the research community in recent years [10, 11, 17, 22, 25, 27, 29, 34]. However, little attention has been paid in the multimedia literature to the on-line nature of the admission control problem for CM database servers. Long and Thakur [23] present simple adversary arguments to show that no on-line algorithm can achieve a *constant* competitive ratio in the context of the Swift distributed I/O architecture. Aggarwal et al. [1, 8] present a competitiveness analysis for a different service model, termed *Shared Video-On-Demand*. Requests are notified of acceptance or rejection within a server-specified time interval (termed *notification interval*) from their arrival. Admitted requests waiting for the same clip, can be batched onto a single stream. They show that allowing for sufficiently large notification intervals (linear in the length of the clips) can guarantee constant competitive ratios for simple scheduling algorithms [1, 8]. The Shared Video-On-Demand model is different from our model of CM service in the sense that it tries to capture the effects of wait tolerance and batching on the number of clients served. Therefore, their results can be viewed as orthogonal to ours. Furthermore, the corresponding analysis assumes that (a) all CM clips have the same length (i.e., time duration) and require the same amount of bandwidth; and (b) any two requests by the same client must be separated by at least the duration of a clip. These assumptions severely limit the applicability of their results to general CM servers.

There is also a significant body of related work in the field of on-line algorithms for bandwidth allocation and circuit routing in communication networks. Lipton and Tomkins [21] study the competitiveness of randomized strategies for the non-preemptive On-line Interval Scheduling (OIS) problem, which essentially corresponds to on-line admission control in a server that can support a *single* playback stream. Under the assumption that the ratio  $\Delta$  of longest to shortest interval is *not known a-priori*, they present an  $O((\log \Delta)^{1+\epsilon})$ -competitive randomized algorithm and show that no  $O(\log \Delta)$ -competitive algorithm can exist. Extensions to their randomized scheme are presented by Faigle et al. [13]. Awerbuch et al. [4] examine the more general problem of non-preemptive circuit routing on tree-structured networks and propose a general randomized technique termed "Classify and Randomly Select". The main idea is to classify on-line events in disjoint classes and then consider only the events that are assigned to a *randomly selected* class. By averaging over all possible random choices, "Classify and Randomly Select" achieves logarithmic competitive ratios (in an expected sense). However, the idea of an admission control scheme that considers only one randomly selected class of user requests and simply ignores all others is obviously not very appealing for CM database servers, since it ignores fundamental requirements such as fairness. Our proposed schemes also employ the idea of on-line classification, but they also are completely *deterministic* with-

out compromising near-optimal competitiveness (for sufficiently large server bandwidth). Awerbuch et al. [3] consider non-preemptive circuit routing on general networks. They present a deterministic scheme (called ROUTE\_OR\_BLOCK) which, assuming that the bandwidth requested by a single circuit never exceeds an  $O(\frac{1}{\log n T_{max}})$  fraction of edge capacity, achieves a competitive ratio of  $O(\log n T_{max})$ , where  $n$  is the number of nodes in the network and  $T_{max}$  is the maximum duration of a call. They also prove that their scheme is near-optimal for deterministic on-line routing. The ROUTE\_OR\_BLOCK algorithm is based on ideas developed for multicommodity network flow problems. Roughly speaking, the main idea is to assign each edge a “length” that is exponential in its current load and route an incoming circuit only if the length of the shortest routing path is less than the “benefit” associated with the circuit. However, as reported by Plotkin [28] and Gawlick [18], ROUTE\_OR\_BLOCK exhibited consistently poor performance in an actual implementation. Ad-hoc changes in the algorithm’s parameters were necessary to improve its behavior. Furthermore, the ROUTE\_OR\_BLOCK scheme itself is rather complex and unintuitive and it is not clear how it can benefit from the knowledge of statistical information, such as request popularities.

Finally, we should note that allowing *preemption* of requests can lead to better competitive ratios for on-line scheduling and admission control problems [7, 9, 14, 15, 20, 33]. However, the assumption of preemptability is unrealistic in the context of CM applications.

### 3 Problem Formulation

We view a CM database server as a “black box” capable of offering a sustained bandwidth capacity of  $B$ . The input sequence consists of a collection of requests  $\bar{\sigma} = \sigma_1, \sigma_2, \dots, \sigma_N$ . The  $i^{th}$  request is represented by the tuple  $\sigma_i = (t_i, l_i, r_i)$ , where  $l_i$ ,  $r_i$  denote the length and bandwidth requirement (respectively) of the requested CM clip and  $t_i$  is the arrival time of  $\sigma_i$ . Given a collection of different requests that are handled by a server (based, for example, on the clips available at the server or the server’s usage patterns), we use  $l_{max}$  ( $l_{min}$ ) to denote the length of the longest (shortest) request. ( $r_{max}$  and  $r_{min}$  are defined similarly.) Finally, we define  $\Delta = \frac{l_{max}}{l_{min}}$  and  $\rho = \frac{r_{max}}{B}$ .

We use *competitive analysis* [30] to measure the performance of different admission control strategies. Our optimization metric is the *total throughput*; that is, the bandwidth-time product over a given sequence of requests. More formally, given an on-line scheduling policy  $A$  and an input sequence  $\bar{\sigma}$ , we define the *benefit* of  $A$  on  $\bar{\sigma}$  as  $V_A(\bar{\sigma}) = \sum_{S_A} l_i \cdot r_i$ , where  $S_A \subseteq \bar{\sigma}$  is the set of requests scheduled by  $A$ . The *competitive ratio* of an on-line algorithm  $A$  is defined as the maximum value  $\kappa(A)$  over all possible request sequences of the ratio of the throughput achieved by the optimal off-line algorithm for a request sequence to the throughput achieved by  $A$  for the same sequence. If  $A$  is a *randomized algorithm*, then the throughput achieved by  $A$  for a request sequence is averaged over all possible “coin flips” of  $A$  [24]. More formally,

$$\kappa(A) = \begin{cases} \sup_{A^*, \bar{\sigma}} \frac{V_{A^*}(\bar{\sigma})}{V_A(\bar{\sigma})} & , \text{ if } A \text{ is deterministic} \\ \sup_{A^*, \bar{\sigma}} \frac{V_{A^*}(\bar{\sigma})}{E_A[V_A(\bar{\sigma})]} & , \text{ if } A \text{ is randomized} \end{cases}$$

where  $\bar{\sigma}$  ranges over all possible request sequences,  $A^*$  ranges

over all off-line (i.e., *clairvoyant*) algorithms, and the expectation  $E_A[\cdot]$  is taken over the random choices of  $A$ . Thus, an algorithm with a *small* competitive ratio is guaranteed to perform close to optimal in all situations. We say that algorithm  $A$  is  $k$ -competitive if  $\kappa(A) \leq k$ .

It is conventional in the analysis of on-line methods to describe things in terms of a game between a player (the on-line algorithm) and an *adversary* (the off-line algorithm), whose goal is to produce a request sequence that would force the player to perform poorly. For randomized algorithms, different models of adversaries have been proposed depending on the adversary’s knowledge of the player’s random choices [24]. The lower bounds presented in this paper assume the “weakest” model of an *oblivious* adversary; that is, an adversary that is oblivious to the random choices made by the on-line algorithm.

## 4 Competitive Analysis of Admission Control

### 4.1 The Greedy/Work-Conserving Policy

The scheduling strategy used as a starting point in our study is the basic *Work-Conserving* ( $WC$ ) scheme traditionally used for admission control in CM servers.  $WC$  is based on the following greedy rule: “schedule request  $\sigma_i$  immediately if the server has at least  $r_i$  bandwidth available at time  $t_i$ ; otherwise, reject  $\sigma_i$ ”. As our results show,  $WC$  offers rather poor performance guarantees in an on-line setting.

First, consider a restricted version of the admission control problem in which all requests require a constant fraction of the server’s bandwidth  $B$ . That is,  $r_i = r$  for all  $i$ . Let  $c = \lfloor \frac{B}{r} \rfloor \geq 1$  denote the number of playback *channels* available at the server<sup>2</sup>. The following theorem establishes the competitiveness of  $WC$  in this setting.

**Theorem 4.1**  $WC$  is  $(1 + \Delta)$ -competitive for scheduling requests with identical bandwidth requirements on-line; that is,  $\kappa(WC) \leq 1 + \Delta$ . Furthermore, this bound is tight. □

In general, a playback request requires an arbitrary portion of the server’s bandwidth  $B$ . This bandwidth requirement depends, for example, on the data encoding method used (e.g., MPEG-1, MPEG-2) or the Quality of Service (QoS) specified by the client. The following theorem shows the effect of this more general model on the competitive factor of the  $WC$  policy.

**Theorem 4.2**  $WC$  is  $\frac{B \cdot (1 + \Delta)}{\max\{B - r_{max}, r_{min}\}}$ -competitive for scheduling requests with different bandwidth requirements on-line. □

Thus, allowing variability along the second dimension (i.e., bandwidth) multiplies the competitiveness of  $WC$  by a factor that depends on  $B$ ,  $r_{max}$ , and  $r_{min}$ . Intuitively, this term captures the effects of the worst-case bandwidth loss due to fragmentation. If  $B \geq r_{max} + r_{min}$ , as will usually be the case for CM servers and requests, then the following corollary applies.

**Corollary 4.1** If  $B \geq r_{max} + r_{min}$  then  $WC$  is  $\frac{1 + \Delta}{1 - \rho}$ -competitive, where  $\rho = \frac{r_{max}}{B}$ . □

<sup>2</sup>Note that, in this case, maximizing throughput is equivalent to maximizing total scheduled request length [21].

Note that the competitiveness bound of  $\frac{1+\Delta}{1-\rho}$  is, in fact, valid regardless of the relative sizes of  $B$  and  $r_{max} + r_{min}$ . Theorem 4.2 just gives a tighter bound when  $B < r_{max} + r_{min}$ . Also, note that for typical CM numbers the fraction  $\rho$  is much smaller than unity. For example, even for the relatively high MPEG-2 rate requirements of 6-8 Mbps, a CM server with a low-end RAID can sustain 40-60 concurrent streams [26, 31]. The denominator in our competitiveness bound agrees with the bounds given by Bar-Noy et al. [7] for the *preemptive* version of the problem. (They avoid dependence on  $\Delta$  through clever use of the preemption mechanism.)

## 4.2 Lower Bounds

In this section, we prove lower bounds on the competitive ratio of any *deterministic or randomized* algorithm for on-line admission control. Our results demonstrate the existence of an exponential separation between the competitive ratio of *WC* and the lower bound on the competitive ratio of any deterministic or randomized algorithm. This clearly suggests the possibility for improvement by using non-greedy schemes. We propose such schemes with near-optimal competitiveness in Section 4.3.

Once again, let us start with the identical bandwidth case (i.e.,  $r_i = r$  for all  $i$ ). A simple adversary argument shows that for the case of a single bandwidth channel (i.e., the OIS problem [21]), there is a lower bound of  $1 + \Delta$  on the competitiveness of any deterministic scheduler. This argument fails when the number of channels is increased. However, as the following theorem shows, no deterministic or randomized admission control scheme can be better than  $\Omega(\log \Delta)$  competitive.

**Theorem 4.3** Any deterministic or randomized on-line admission control algorithm for CM requests with identical bandwidth requirements has a competitive ratio of  $\Omega(\log \Delta)$ .  $\square$

Similar lower bounds on the competitive ratio hold for the variable bandwidth case. Again, the effect of bandwidth fragmentation introduces a multiplicative factor of  $\frac{1}{1-\rho}$ .

**Theorem 4.4** Consider a sequence of CM requests with variable bandwidth requirements. Then:

- (1) Any *deterministic* on-line admission control algorithm has a competitive ratio of (a)  $\Omega(\frac{\rho \cdot (1+\Delta)}{1-\rho})$ , if  $\rho > \frac{1}{2}$ ; and, (b)  $\Omega(\frac{\log \Delta}{1-\rho})$ , otherwise.
- (2) Any *randomized* on-line admission control algorithm has a competitive ratio of  $\Omega(\frac{\log \Delta}{1-\rho})$ , if  $\rho \leq \frac{1}{2 \log \Delta}$ .  $\square$

We should note that Awerbuch et al. [3] also proved an  $\Omega(\log \Delta)$  lower bound for *deterministic* on-line circuit routing in the case of requests with identical bandwidth requirements. However, our lower bounds for the more general variable bandwidth case also demonstrate the effect of the maximum bandwidth demand ( $\rho$ ) which was not factored into their results. Furthermore, we have shown that the logarithmic lower bounds cannot be improved upon through the use of randomization.

## 4.3 Bandwidth Prepartitioning Policies

We now propose novel deterministic admission control policies that guarantee near-optimal competitive ratios for reasonably large bandwidth capacities. Our policies are based on *prepartitioning* the bandwidth capacity of a CM server among requests of different length. Roughly speaking, the basic idea of the bandwidth prepartitioning schemes is to isolate requests with large differences in length, thus ensuring that short requests cannot “steal” the entire server bandwidth from longer (and, more profitable) requests<sup>3</sup>.

The first policy we introduce is termed *Simple Bandwidth Prepartitioning (SBP)* and is depicted in Figure 1<sup>4</sup>. The *SBP* algorithm exploits the server’s knowledge of the  $\Delta$  ratio by classifying requests to channel groups based on their length and then using a *WC* policy within each group. The idea is that by classifying the requests into different partitions according to their length range, we are ensuring that the maximum to minimum length ratio is bounded by a constant *within each partition*.

### Simple Bandwidth Prepartitioning

1. Divide the available bandwidth  $B$  into  $\lceil \log \Delta \rceil$  partitions  $B_1, \dots, B_{\lceil \log \Delta \rceil}$ , where the size of the  $i^{th}$  partition is  $|B_i| = \frac{B}{\lceil \log \Delta \rceil}$ .
2. For each arriving request  $\sigma_j = (t_j, l_j, r_j)$ 
  - 2.1 Let  $i \in \{1, \dots, \lceil \log \Delta \rceil\}$  be such that:  $2^{i-1} \cdot l_{min} \leq l_j < 2^i \cdot l_{min}$  (allowing for  $l_j = 2^i \cdot l_{min}$  if  $i = \lceil \log \Delta \rceil$ ).
  - 2.2 If the amount of free bandwidth in partition  $B_i$  is less than  $r_j$ , then reject  $\sigma_j$ ; Otherwise, schedule  $\sigma_j$  in partition  $B_i$ ;

Figure 1: Algorithm *SBP*.

The following theorem shows that this simple prepartitioning scheme results in a significant improvement in the competitive ratio for CM servers with bandwidth  $B$  larger than  $r_{max} \cdot \lceil \log \Delta \rceil$ . This requirement is typically satisfied by today’s servers, even for large values of  $r_{max}$  and  $\Delta$ . For example, if  $r_{max} = 8$  Mbps and  $l_{max} = 120 \cdot l_{min}$ , then  $r_{max} \cdot \lceil \log \Delta \rceil = 56$  Mbps, i.e., less than the transfer rate of a *single* high-end magnetic disk [27].

**Theorem 4.5** Assume  $\rho = \frac{r_{max}}{B} < \frac{1}{\lceil \log \Delta \rceil}$  (or, equivalently,  $c > \lceil \log \Delta \rceil$  for the identical bandwidth case). Then the *SBP* admission control policy is:

- (1)  $3 \cdot \lceil \log \Delta \rceil$ -competitive for the identical bandwidth case; and,
- (2)  $\frac{3 \cdot \lceil \log \Delta \rceil}{1-\rho \cdot \lceil \log \Delta \rceil}$ -competitive for the variable bandwidth case.  $\square$

<sup>3</sup>Note that, if  $\Delta = 1$  (i.e., all requests have identical lengths) then simple *WC* offers optimal competitiveness. Thus, we will assume that  $\Delta > 1$  or, equivalently,  $\log \Delta > 0$  in the remainder of this paper.

<sup>4</sup>We describe our policies in terms of the more general variable bandwidth case. The restriction to identical bandwidth requests should be straightforward.

Thus, by merely isolating different length ranges, the *SBP* admission control policy improves the competitiveness of *WC* from linear to logarithmic in  $\Delta$ , at least for the identical bandwidth case. The main idea behind *SBP* is that in order to be competitive under a worst-case scenario, the scheduler should not allow short duration requests to monopolize the server's bandwidth. However, *SBP* can also suffer from bandwidth fragmentation in the variable bandwidth case. In the worst case, bandwidth approximately equal to  $r_{max}$  is lost *in each partition*, leading to a total bandwidth loss of  $r_{max} \cdot \lceil \log \Delta \rceil$  in the server. Intuitively, we would like to be able to “combine” these bandwidth fragments to allow for incoming requests to be scheduled across partitions, especially if these requests are long since this implies more guaranteed profit.

The *Down-shift Bandwidth Prepartitioning (DBP)* policy depicted in Figure 2 is based exactly on these observations. As in *SBP*, the *DBP* algorithm also prohibits short requests from monopolizing the server, *but* it also allows longer (and thus, more profitable requests) to be “down-shifted” to lower groups and steal unused bandwidth that would otherwise be dedicated to shorter requests. Theorem 4.6 shows that incorporating this change does not compromise logarithmic competitiveness.

**Theorem 4.6** Assume  $\rho = \frac{r_{max}}{B} < \frac{1}{\lceil \log \Delta \rceil}$  (or, equivalently,  $c > \lceil \log \Delta \rceil$  for the identical bandwidth case). Then the *DBP* admission control policy is:

- (1)  $(1+8 \cdot \lceil \log \Delta \rceil)$ -competitive for the identical bandwidth case; and,
- (2)  $\left(1 + \frac{8 \cdot \lceil \log \Delta \rceil}{1-\rho \cdot \lceil \log \Delta \rceil}\right)$ -competitive for the variable bandwidth case.

□

### Down-shift Bandwidth Prepartitioning

1. Divide the available bandwidth  $B$  into  $\lceil \log \Delta \rceil$  partitions  $B_1, \dots, B_{\lceil \log \Delta \rceil}$ , where the size of the  $i^{th}$  partition is  $|B_i| = \frac{B}{\lceil \log \Delta \rceil}$ .
2. For each arriving request  $\sigma_j = (t_j, l_j, r_j)$ 
  - 2.1 Let  $i \in \{1, \dots, \lceil \log \Delta \rceil\}$  be such that:  $2^{i-1} \cdot l_{min} \leq l_j < 2^i \cdot l_{min}$  (allowing for  $l_j = 2^i \cdot l_{min}$  if  $i = \lceil \log \Delta \rceil$ ).
  - 2.2 If the total amount of free bandwidth in  $B_1 \cup \dots \cup B_i$  is less than  $r_j$ , then reject  $\sigma_j$ ; Otherwise, schedule  $\sigma_j$  using available bandwidth from partitions  $B_i, B_{i-1}, \dots, B_1$  in that order.

Figure 2: Algorithm *DBP*.

Corollary 4.2 follows directly from Theorems 4.5 and 4.6. Combined with the lower bounds in Section 4.2, Corollary 4.2 establishes the *near-optimality* of the *SBP* and *DBP* policies for the variable bandwidth case, assuming that the server bandwidth  $B$  is larger than  $2 \cdot r_{max} \cdot \lceil \log \Delta \rceil$ . Again, this is a requirement that is typically satisfied by today's CM

servers and applications. (See the discussion before Theorem 4.5.) Note that even smaller competitive ratios can be obtained if  $B > k \cdot r_{max} \cdot \lceil \log \Delta \rceil$ , where  $k > 2$ .

**Corollary 4.2** Assume  $\rho = \frac{r_{max}}{B} < \frac{1}{2 \cdot \lceil \log \Delta \rceil}$ . Then:

- (1) The *SBP* admission control policy is  $6 \cdot \lceil \log \Delta \rceil$ -competitive for the variable bandwidth case; and,
- (2) The *DBP* admission control policy is  $(1 + 16 \cdot \lceil \log \Delta \rceil)$ -competitive for the variable bandwidth case.

□

Although the constants in the competitiveness bounds we have shown for *DBP* are larger than those of *SBP*, we conjecture that they can be improved. To support our conjecture, note that for the identical bandwidth case when a request of length  $l_j \in [2^{i-1} \cdot l_{min}, 2^i \cdot l_{min})$  is rejected in *SBP*, the scheduler can guarantee that the benefit of running requests is at least  $\frac{c}{\lceil \log \Delta \rceil} \cdot 2^{i-1} \cdot l_{min}$ , whereas with *DBP* the corresponding guaranteed benefit is at least:

$$\frac{c}{\lceil \log \Delta \rceil} \cdot \sum_{k=1}^i 2^{k-1} \cdot l_{min} = \frac{c}{\lceil \log \Delta \rceil} \cdot (2^i - 1) \cdot l_{min},$$

that is, nearly double the benefit for *SBP*. Of course, the main advantage of *DBP* over *SBP* is that, by “down-shifting”, it can significantly reduce the effects of bandwidth fragmentation in the variable bandwidth case. A formal proof of improved competitive ratios for *DBP* is left as an open problem for future research.

Even though *SBP* and *DBP* guarantee logarithmic competitiveness under a worst-case scenario, they may also severely underutilize the server in average cases. For example, when all the requests address the shortest group of clips residing on the server, both schemes will end up utilizing only  $\frac{1}{\lceil \log \Delta \rceil}$  of the available bandwidth. This is clearly undesirable. We now propose a novel on-line admission control policy that employs the intuition of partitioning schemes (to avoid worst-case scenarios for *WC*) within a framework that also allows for good average-case performance. Roughly speaking, the idea is to use the methodology of *DBP* but define the sizes of the bandwidth partitions  $B_i$  as a function of the popularities and/or the lengths of all requests in the length range  $[2^{i-1} \cdot l_{min}, 2^i \cdot l_{min})$ . The resulting admission control scheme, termed *Popularity-based Bandwidth Prepartitioning (PBP)*, is depicted in Figure 3. Note that *PBP* is given in parameterized form with the parameter  $f$  being the specific function of popularities and lengths used to define the partition sizes. In Section 5, we describe two specific choices for  $f$  used in our preliminary experimental study. The *PBP* admission control scheme relies on the assumption that request (i.e., clip) popularities can be estimated with reasonable accuracy (e.g., using a “moving window” prediction method [22]). Clearly, taking popularities into account is necessary to avoid worst-case scenarios for *DBP* (i.e., when the most frequent requests are also the shortest). In fact, assuming that requests are independent and the given popularities are accurate, we can give simple arguments based on *Chernoff bounds* [24] to show that the probability of a worst-case “loss” for *PBP* (with specific choices for  $f$ ) is exponentially small. The interested reader is referred to the full version of this paper [16].

### Popularity-based Bandwidth Prepartitioning[ $f$ ]

1. Let  $p_j$  denote the probability that the length of an incoming request is  $l_j$  (i.e., the popularity of  $l_j$ ). Let  $PL_i$  denote the set of (popularity, length) pairs with lengths in the  $i^{\text{th}}$  range; that is,

$$PL_i = \{ (p_j, l_j) \mid l_j \in [2^{i-1} \cdot l_{min}, 2^i \cdot l_{min}) \}.$$

2. Divide the available bandwidth  $B$  into  $\lceil \log \Delta \rceil$  partitions  $B_1, \dots, B_{\lceil \log \Delta \rceil}$ , with

$$|B_i| = \frac{f(PL_i)}{\sum_i f(PL_i)} \cdot B.$$

3. For each arriving request  $\sigma_j = (t_j, l_j, r_j)$ 
  - 3.1 Let  $i \in \{1, \dots, \lceil \log \Delta \rceil\}$  be such that:  $2^{i-1} \cdot l_{min} \leq l_j < 2^i \cdot l_{min}$  (allowing for  $l_j = 2^i \cdot l_{min}$  if  $i = \lceil \log \Delta \rceil$ ).
  - 3.2 If the total amount of free bandwidth in  $B_1 \cup \dots \cup B_i$  is less than  $r_j$ , then reject  $\sigma_j$ ; Otherwise, schedule  $\sigma_j$  using available bandwidth from partitions  $B_i, B_{i-1}, \dots, B_1$  in that order.

Figure 3: Algorithm  $\mathcal{PBP}$

## 5 Experimental Study

In this section, we describe the results of a preliminary set of experiments we have conducted with the  $\mathcal{WC}$  and  $\mathcal{PBP}$  strategies for on-line admission control. Since our competitiveness results clearly demonstrate the superiority of prepartitioning schemes with respect to worst-case scenarios, our goal was to ensure that the worst-case guarantees did not impair *average-case performance*. We start by presenting our experimental testbed and methodology.

### 5.1 Experimental Testbed

To examine the average-case behavior of the  $\mathcal{WC}$  and  $\mathcal{PBP}$  schemes, we have experimented with three distinct random arrival patterns:

- **Poisson Arrivals.** Requests of different lengths arrive at the server according to a Poisson process model with an arrival rate of  $\lambda$ . This is a plausible probabilistic model for servers with a reasonably steady traffic flow (e.g., video servers in scientific research labs serving clips of recorded experiments to scientists around the globe).
- **Bursty Arrivals.** Requests of different lengths arrive at the server in bursts at regular intervals of time (termed *burst separation*). Each such burst itself consists of a sequence of *request batches*, where each batch consists of requests of identical length arriving during a very short period of time. The batch arrivals are again modeled as a Poisson process with an arrival rate of  $\lambda$ . This workload is intended to model “rush-hour traffic” situations in CM servers.

- **Poisson + Short Burst Arrivals.** Long requests arrive at the server according to a Poisson process model with an arrival rate of  $\lambda_{long}$ . At the same time, bursts of short requests arrive based on a Poisson process with an arrival rate of  $\lambda_{short}$ . This workload model combines some features of the previous two models. It is intended to represent situations where servers operating under a relatively steady flow of long requests (e.g., movies or sports events), occasionally have to handle bursts of short requests (e.g., the 6 o'clock news).

In most of our experiments, the request lengths were sampled from a discrete set of values between 5 and 150 minutes, with sampling probabilities (i.e., popularities) taken from a Zipfian distribution with skew parameter  $z$  [35]. We varied this skew parameter from 0.0 (uniform) to 2.0 (very skewed). Results were obtained for three different models of correlation between request lengths and popularities:

- **Positive.** Larger popularities are assigned to longer requests.
- **Negative.** Larger popularities are assigned to shorter requests.
- **Random.** No length/popularity correlation exists; that is, the values of the Zipfian probability vector are assigned to the different request lengths in a random manner.

We also experimented with two different choices for the  $f$  function parameter of the  $\mathcal{PBP}$  scheme. The first choice  $f_1$  captured the cumulative popularity of a length range, that is  $f_1(PL_i) = \sum_{(p,l) \in PL_i} p$ . The second choice  $f_2$  was the total popularity-length product of a range, that is  $f_2(PL_i) = \sum_{(p,l) \in PL_i} p \cdot l$ .

In our experiments for the identical bandwidth case, we assumed a server with 100 available channels. For the variable bandwidth case, we varied the server’s sustained bandwidth capacity between 100 and 250 Megabits per second (Mbps) and selected the rate requirement of a request randomly between 500 Kbps and 8 Mbps. The parameter values are summarized in Table 1.

System Parameter	Value
Server Bandwidth Capacity	100 channels / 100–250 Mbps
Request Lengths	5 minutes – 150 minutes
Request Rates (Variable Bandwidth Case)	500 Kbps - 8 Mbps
Zipfian Popularity Skew ( $z$ )	0.0 - 2.0

Table 1: Experiment Parameter Settings

For each different combination of input parameters, we modeled the system behavior under each scheduling policy for 20,000 minutes of simulated time and 10 randomly generated request sequences. The results presented here represent the averages over these 10 runs of the system. In all cases, the comparison metric was the *fraction of the server capacity effectively utilized*; that is, the ratio

$$\frac{V_A(\bar{\sigma})}{(\text{server bandwidth}) \times (\text{simulation time})},$$

for each scheduler  $A$  and request sequence  $\bar{\sigma}$ .

## 5.2 Experimental Results

We now present an overview of our experimental comparison of the  $\mathcal{WC}$  and  $\mathcal{PBP}$  schemes for the (general) variable bandwidth case. Similar results were obtained for identical bandwidth requests. For the numbers presented here, the request lengths were sampled (based on  $z$  and the model of correlation) from the collection  $\{5, 10, 15, 90, 120, 150\}$  (in minutes), the request rates were selected (uniformly) from the set  $\{0.5, 1.5, 3.0, 4.5, 6.0, 8.0\}$  (in Mbps), and the server bandwidth capacity was 250 Mbps. The plots shown in this section are indicative of the results obtained for other values of request and server parameters.

We focus our discussion on  $\mathcal{PBP}[f_2]$ ; that is,  $\mathcal{PBP}$  using the “total popularity $\times$ length” partitioning criterion, since it exhibited uniformly better performance than  $\mathcal{PBP}[f_1]$  in our experiments. We should stress, however, that even with “cumulative popularity” partitioning,  $\mathcal{PBP}$  outperformed  $\mathcal{WC}$  by a significant margin for our “bursty” workloads.

The first set of experiments studied the relative effectiveness of the  $\mathcal{WC}$  and  $\mathcal{PBP}$  schemes under Poisson arrivals for different values of the Zipfian skew parameter  $z$  and different length/popularity correlations. Figure 4(a) shows the performance of the schemes as a function of the Poisson arrival rate  $\lambda$  for  $z = 0.6$  and random length/popularity correlation. Our basic finding is that, by exploiting its knowledge of clip popularities  $\mathcal{PBP}$  is able to do at least as good as  $\mathcal{WC}$  in all cases. We should mention that we also experimented with different models of the arrival process (e.g., using uniformly rather than exponentially distributed inter-arrival times) that also led to the same conclusions regarding the relative performance of the strategies under random arrivals.

The second set of experiments concentrated on the relative performance of the algorithms under the Bursty Arrival model described in the previous section. We studied the server utilization as a function of the length of the burst separation interval as well as the size of a batch of arrivals for different values of the  $z$  and  $\lambda$  parameters, the size (i.e., number of batches) of a burst, and different modes of correlation. Figure 4(b) shows the server utilization as a function of the burst separation interval for batch size equal to 40,  $z = 0.6$ , batch arrival rate  $\lambda = 0.8$ , burst size equal to 10, and random length/popularity correlation. (We show burst separations decreasing from left to right as this reflects increasing load, as in Figure 4(a).) Our results show that under such conditions,  $\mathcal{PBP}$  outperforms  $\mathcal{WC}$  by an average margin of 15% - 40%. Note that the “jump” observed in the curves as the burst separation approaches 150 minutes is caused by our specific choice of request lengths and our model of “bursty” arrivals. The numbers from the same experiment but for *positive* length/popularity correlation (i.e., longer requests are more popular) are depicted in Figure 5(a).  $\mathcal{WC}$  clearly performs better under the positive correlation assumption, since it is able to allocate more of its channels to the more popular (and, more profitable) long requests. Still,  $\mathcal{PBP}$  continues to outperform  $\mathcal{WC}$  by up to 25%. Figure 5(b) shows the results of the same experiment but for *negative* length/popularity correlation (i.e., shorter requests are more popular). Under such scenarios, our results show that the relative improvement offered by  $\mathcal{PBP}$  over  $\mathcal{WC}$  can reach 50% - 60%. A different perspective is depicted in Figure 6(a), where server utilization (for the same parameter values and negative correlation) is given as a function of the batch size for a fixed burst separation of 180 minutes. Note that as the batch size increases, the Bursty Arrival model gives rise to worst-case scenarios for

$\mathcal{WC}$ , where a large batch of short requests can flood the server leaving no capacity for a following batch of larger requests. On the other hand, our  $\mathcal{PBP}$  scheme is capable of maintaining a reasonable level of utilization under all circumstances.

The final set of experiments studied the behavior of the algorithms under the combined Poisson + Short Bursts arrival process. We concentrated on a particular scenario which, we believe, is common in Video-On-Demand environments. Specifically, we assumed that the server is working close to capacity serving requests for long (i.e.,  $\{90, 120, 150\}$  minutes) movies but occasionally has to handle bursts of short (i.e.,  $\{5, 10, 15\}$  minutes) requests. That is,  $\lambda_{long}$  was selected large enough to ensure high system utilization and we studied the server utilization as a function of  $\lambda_{short}$ . All length popularities were assumed uniform for this experiment. The results depicted in Figure 6(b) show that, under this scenario,  $\mathcal{PBP}$  can offer a 10% - 15% performance improvement over  $\mathcal{WC}$ , even at high levels of system utilization.

## 6 Conclusions and Future Work

In this paper, we have addressed the admission control problem associated with CM database servers from a novel, on-line perspective. Using server throughput as our optimization metric, we showed that the traditionally used Work-Conserving policy has a competitive ratio of  $\frac{1+\Delta}{1-\rho}$ , where  $\Delta$  is the ratio of the maximum to minimum request length and  $\rho$  is the maximum fraction of the server’s bandwidth that a request can demand. We developed novel admission control strategies based on the simple idea of prepartitioning the bandwidth capacity of the server among requests of different length and proved that that our strategies are  $O(\log \Delta)$ -competitive for sufficiently large server bandwidth. We also showed an  $\Omega(\frac{\log \Delta}{1-\rho})$  lower bound on the competitive ratio of any deterministic *or* randomized algorithm for the problem, thus establishing that our bandwidth prepartitioning algorithms are within a multiplicative constant of the optimal on-line strategy. Based on the intuition gained from our competitiveness results, we proposed prepartitioning schemes that make use of request popularities to ensure good average-case as well as robust worst-case performance, and experimentally verified their effectiveness against the Work-Conserving policy.

We believe that the analytical and experimental results presented in this paper offer new insights to other optimization problems that arise in CM data management. For example, consider the problem of data placement and static load balancing in distributed CM servers. Briefly, the problem can be described as follows: Given a collection of continuous media clips with lengths ( $l_i$ ), rates ( $r_i$ ), and expected popularity (or, probability of access,  $p_i$ ), determine a “good quality” mapping of these clips to a collection of servers, where each server is characterized by a bandwidth capacity ( $B_j$ ) and a storage capacity ( $S_j$ ) and a clip can be mapped to more than one servers (i.e., replication of clips is allowed).

Traditionally, the goal of data placement schemes in this setting is to balance the expected bandwidth load (according to the popularities  $\{p_i\}$ ) across the available servers under the given server storage constraints [11, 22, 34]. This model of “popularity-based data placement” aims at achieving good system utilization and balanced system load in an average sense. On the other hand, our competitiveness results indicate that to ensure *robust* system performance, a placement strategy should also try to achieve some sec-

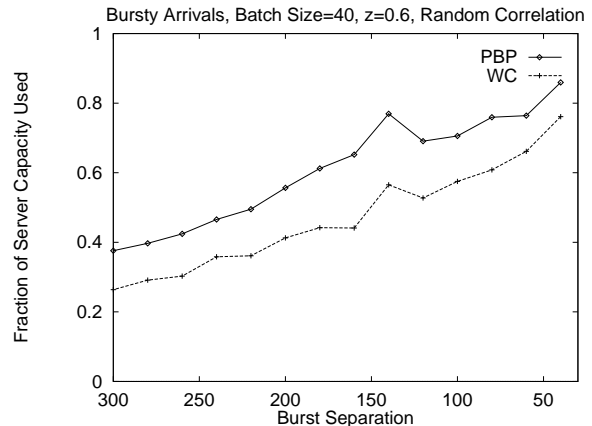
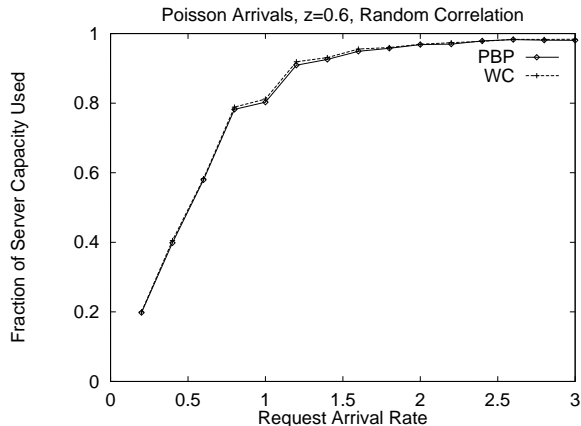


Figure 4: (a) Server throughput under Poisson arrivals. (b) Server throughput under bursty arrivals (random correlation).


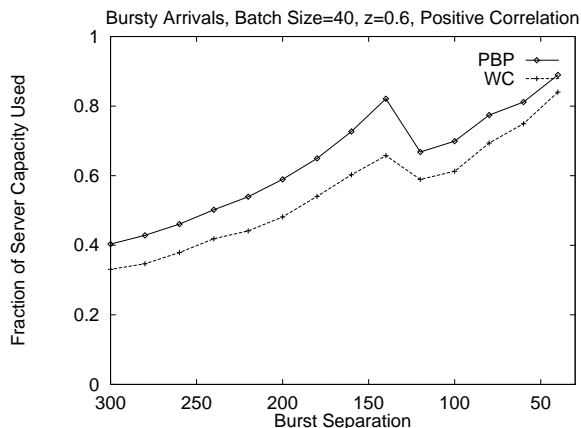


Figure 5: (a) Server throughput under bursty arrivals (positive correlation). (b) Server throughput under bursty arrivals (negative correlation).

ondary “goals”. One such goal, for example, would be to place clips with large bandwidth requirements on servers with large bandwidth capacities to guarantee small  $\rho$  fractions for each server. As another example, the placement policy should try to replicate short clips across many servers, so that there are many possibilities of dynamically re-assigning (e.g., using a baton passing primitive [34]) streams delivering these clips to different servers. This obviously reduces the probability of a short request causing the rejection of a long request from the system, which is the worst-case scenario in all our competitiveness results. Achieving such secondary data placement goals is especially important in order to ensure good system utilization under short-term fluctuations of the load away from the averages  $\{p_i\}$  or overload situations where some client requests simply must be rejected. A detailed investigation of the problem is left for future research.

The competitive analysis framework and results presented in this paper suggest several directions for future work. First, this work has only considered the bandwidth resource. Considering servers with multiple scarce resources poses an interesting challenge. For example, a request may also need a given amount of memory at the server in order to meet its performance requirements. This memory requirement can

be either specified by the request itself (e.g., leaky bucket regulated traffic) or assigned by the server to meet the request’s performance goals. Given the limited amount of server memory, the admission control mechanism needs to consider both the memory and the bandwidth requirements of a request. Second, we plan to use our results and methodology as a starting step towards a formal study of dynamic on-line load balancing in multimedia storage servers. Prior work [11, 34] has explored the use of techniques such as dynamic stream re-assignments and dynamic replica creation/coalescing only within ad-hoc schemes, without providing any strong performance guarantees or exploiting the wealth of theoretical results on on-line load balancing (see, for example, [2, 5, 6, 32]). Finally, incorporating the concept of *equivalent bandwidth* into our on-line analysis is another interesting problem. Equivalent bandwidth is fundamental for providing sessions with statistical performance guarantees. Since the equivalent bandwidth of a collection of sessions is typically a simple function of each session’s equivalent bandwidth, this concept simplifies the admission control for applications with statistical performance guarantees [12, 19].



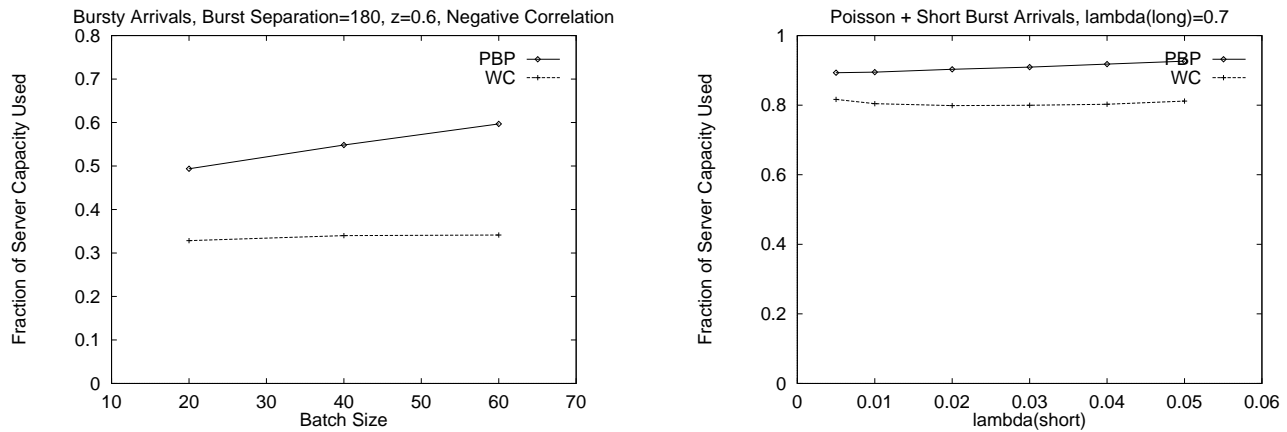


Figure 6: (a) Server throughput under bursty arrivals as a function of batch size (negative correlation). (b) Server throughput under Poisson+Short Bursts.

## References

- [1] Sudhanshu Aggarwal, Juan A. Garay, and Amir Herzberg. "Adaptive Video On Demand". In *Proceedings of the 3rd Annual European Symposium on Algorithms (ESA '95)*, Corfu, Greece, September 1995.
- [2] James Aspnes, Yossi Azar, Amos Fiat, Serge Plotkin, and Orli Waarts. "On-Line Load Balancing with Applications to Machine Scheduling and Virtual Circuit Routing". In *Proceedings of the 25th Annual ACM Symposium on the Theory of Computing*, pages 623–631, San Diego, California, May 1993.
- [3] Baruch Awerbuch, Yossi Azar, and Serge Plotkin. "Throughput-Competitive On-Line Routing". In *Proceedings of the 34th Annual IEEE Symposium on Foundations of Computer Science*, pages 32–40, Palo Alto, California, November 1993.
- [4] Baruch Awerbuch, Yair Bartal, Amos Fiat, and Adi Rosén. "Competitive Non-Preemptive Call Control". In *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 312–320, Arlington, Virginia, January 1994.
- [5] Yossi Azar. "On-line Load Balancing". Unpublished Manuscript (available from <http://www.math.tau.ac.il/~azar/>), 1997.
- [6] Yossi Azar, Andrei Z. Broder, and Anna R. Karlin. "On-line Load Balancing". *Theoretical Computer Science*, 130:73–84, 1994.
- [7] Amotz Bar-Noy, Ran Canetti, Shay Kutten, Yishay Mansour, and Baruch Schieber. "Bandwidth Allocation with Preemption". In *Proceedings of the 27th Annual ACM Symposium on the Theory of Computing*, pages 616–625, Las Vegas, Nevada, May 1995.
- [8] Amotz Bar-Noy, Juan A. Garay, Amir Herzberg, and Sudhanshu Aggarwal. "Sharing Video on Demand". Unpublished Manuscript (available from <http://www.eng.tau.ac.il/~amotz/publications.html>), 1996.
- [9] S. Baruah, G. Koren, D. Mao, B. Mishra, A. Raghunathan, L. Rosier, D. Shasha, and F. Wang. "On the Competitiveness of On-Line Real-Time Task Scheduling". In *Proceedings of the 12th IEEE Real-Time Systems Symposium*, pages 106–115, San Antonio, Texas, December 1991.
- [10] Mon-Song Chen, Dilip D. Kandlur, and Philip S. Yu. "Optimization of the Grouped Sweeping Scheduling (GSS) with Heterogeneous Multimedia Streams". In *Proceedings of ACM Multimedia '93*, pages 235–242, Anaheim, California, August 1993.
- [11] Asit Dan and Dinkar Sitaram. "An Online Video Placement Policy based on Bandwidth to Space Ratio (BSR)". In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data*, San Jose, California, May 1995.
- [12] A. Elwalid, D. Mitra, and R. H. Wentworth. "A New Approach for Allocating Buffers and Bandwidth to Heterogeneous, Regulated Traffic in an ATM Node". *IEEE Journal on Selected Areas in Communication*, 13(6):1115–1127, August 1995.
- [13] U. Faigle, R. Garbe, and W. Kern. "Randomized Online Algorithms for Maximizing Busy Time Interval Scheduling". *Computing*, 56:95–104, 1996.
- [14] Ulrich Faigle and Willem M. Nawijn. "Note on Scheduling Intervals On-line". *Discrete Applied Mathematics*, 58:13–17, 1995.
- [15] Juan A. Garay, Inder S. Gopal, Shay Kutten, Yishay Mansour, and Moti Yung. "Efficient On-Line Call Control Algorithms". *Journal of Algorithms*, 23:180–194, 1997.
- [16] Minos N. Garofalakis, Yannis E. Ioannidis, Banu Özden, and Avi Silberschatz. "Throughput - Competitive Admission Control for Continuous Media Databases". Tech. Memorandum BL0112300-971126-35TM, Bell Laboratories, November 1997.
- [17] Minos N. Garofalakis, Banu Özden, and Avi Silberschatz. "Resource Scheduling in Enhanced Pay-Per-View Continuous Media Databases". In *Proceedings of*

*the 23rd International Conference on Very Large Data Bases*, pages 516–525, Athens, Greece, August 1997.

- [18] Rainer Gawlick. “Admission Control and Routing: Theory and Practice”. PhD thesis, Laboratory for Computer Science, Massachusetts Institute of Technology, June 1995. (Technical Report MIT/LCS/TR-679).
- [19] R. Guerin, H. Ahmadi, and M. Nagshineh. “Equivalent Capacity and its Application to Bandwidth in High Speed Networks”. *IEEE Journal on Selected Areas in Communication*, 9(7):968–981, September 1991.
- [20] Gilad Koren and Dennis Shasha. “An Optimal Scheduling Algorithm with a Competitive Factor for Real-Time Systems”. Technical Report TR 572, Department of Computer Science, New York University, July 1991.
- [21] Richard J. Lipton and Andrew Tomkins. “Online Interval Scheduling”. In *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 302–311, Arlington, Virginia, January 1994.
- [22] T.D.C. Little and D. Venkatesh. “Popularity-Based Assignment of Movies to Storage Devices in a Video-on-Demand System”. *ACM Multimedia Systems*, 2:280–287, 1995.
- [23] Darrell D.E. Long and Madhukar N. Thakur. “Scheduling Real-Time Disk Transfers for Continuous Media Applications”. In *Proceedings of the Twelfth IEEE Symposium on Mass Storage Systems*, pages 227–232, 1993.
- [24] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [25] Banu Özden, Rajeev Rastogi, and Avi Silberschatz. “A Framework for the Storage and Retrieval of Continuous Media Data”. In *Proceedings of the 1995 International Conference on Multimedia Computing and Systems*, pages 2–13, Washington, D.C., May 1995.
- [26] Banu Özden, Rajeev Rastogi, and Avi Silberschatz. “Disk Striping in Video Server Environments”. *IEEE Data Engineering Bulletin*, 18(4):4–16, December 1995. (Special Issue on Multimedia Information Systems).
- [27] Banu Özden, Rajeev Rastogi, and Avi Silberschatz. “The Storage and Retrieval of Continuous Media Data”. In *“Multimedia Database Systems: Issues and Research Directions”, V.S. Subrahmanian and Sushil Jajodia (Eds.)*, pages 237–261. Springer-Verlag, 1996.
- [28] Serge Plotkin. “Competitive Routing of Virtual Circuits in ATM Networks”. *IEEE Journal on Selected Areas in Communications*, pages 1128–1136, August 1995. (Invited Paper).
- [29] P. Venkat Rangan and Harrick M. Vin. “Designing File Systems for Digital Video and Audio”. In *Proceedings of the Thirteenth ACM Symposium on Operating Systems Principles*, pages 81–94, Monterey, California, October 1991.
- [30] Daniel D. Sleator and Robert E. Tarjan. “Amortized Efficiency of List Update and Paging Rules”. *Communications of the ACM*, 28(2):202–208, February 1985.
- [31] Yuewei Wang, Jonathan C.L. Liu, David H.C. Du, and Jenwei Hsieh. “Video File Allocation over Disk Arrays for Video-on-Demand”. In *Proceedings of the 1996 International Conference on Multimedia Computing and Systems*, pages 160–163, Hiroshima, Japan, June 1996.
- [32] Jeffery Westbrook. “Load Balancing for Response Time”. In *Proceedings of the 3rd Annual European Symposium on Algorithms (ESA ’95)*, pages 355–368, Corfu, Greece, September 1995.
- [33] Gerhard J. Woeginger. “On-line Scheduling of Jobs with Fixed Start and End Times”. *Theoretical Computer Science*, 130:5–16, 1994.
- [34] Joel L. Wolf, Philip S. Yu, and Hadas Shachnai. “DASD Dancing: A Disk Load Balancing Optimization Scheme for Video-on-Demand Computer Systems”. In *Proceedings of the 1995 ACM SIGMETRICS Conference on Measurement & Modeling of Computer Systems*, Ottawa, Canada, May 1995.
- [35] George Kingsley Zipf. *Human Behavior and the Principle of Least Effort – An Introduction to Human Ecology*. Addison-Wesley Press, Inc., Cambridge 42, Massachusetts, 1949.