

Techniques for improving the throughput of VBR streams *

A. L. Narasimha Reddy
Ravi Wijayaratne¹

Dept. of Elec. Engg.
¹Dept. of Comp. Sci.
Texas A & M University
214 Zachry
College Station, TX 77843-3128
{reddy,ravi}@ee.tamu.edu

Abstract

In this paper, we propose a novel mechanism for providing deterministic service for Variable Bit Rate (VBR) streams at the disk. Previous approaches have relied on peak-rate of the stream for admission control for providing deterministic service. The proposed scheme allows statistical multiplexing of VBR streams at the storage system while providing deterministic service guarantees. We show that the proposed scheme can significantly improve stream throughput compared to peak-rate based schemes. We also evaluate the impact of other strategies such as data smoothing, statistical guarantees and higher stream startup latencies for improving the throughput of VBR streams. We show that stream startup latency can be effectively traded off for improving the stream throughput. We also show that smoothing and statistical guarantees do not provide significant additional improvements in stream throughput beyond the proposed approach.

1 Introduction

System level support of continuous media has been receiving wide attention. Continuous media impose timing requirements on the retrieval and delivery of data unlike traditional data such as text and images. Timely retrieval and delivery of data requires that the system and network pay attention to notions of time and deadlines. Data retrieval is handled by the I/O system (File system, disk drivers, disks etc.) and the delivery is handled by the network system (network software and the network). In this paper, we will look at the data retrieval problem.

Different levels of service can be provided for continuous media. *Deterministic service* provides guarantees that the required data will be retrieved in time. *Statistical service* provides statistical guarantees about

data retrieval, e.g., 99.9% of the requested blocks will be retrieved in time. Data streams can be classified as Constant Bit Rate (CBR) or Variable Bit Rate (VBR) depending on whether the stream requests the same amount of data in an interval. Disk service tends to incur random service time costs because of the random seek and latency overheads involved in disk transfers. Much work has been done on providing deterministic service for CBR streams [1, 2, 3, 4] and many projects have started building servers for such service [5, 6, 7, 8]. Recently, statistical service for VBR streams received attention [9, 10]. In this paper, we address the problem of providing deterministic service for VBR streams.

Providing deterministic service at disk is complicated by the random seek and latency overheads involved in a disk transfer. This problem has been addressed effectively by suitable disk scheduling policies [11, 12]. These scheduling policies group a number of requests into rounds or batches and service the requests in a round using a disk seek optimizing policy such as SCAN. Then the service time for the entire round can be bounded to provide guarantees. This strategy works well with CBR streams. However, with VBR streams, the workload changes from round to round and hence such an approach will have to consider the variations in load for providing guarantees. Long term behavior of MPEG video streams exhibit normal distribution [9, 13]. This observation can be used to compute the statistical probabilities of workload variations when a number of streams require service at an I/O system. Based on these probabilities, statistical service can be provided for VBR streams.

How to provide deterministic service for VBR streams? This problem has not received much attention and is the focus of this paper. Providing deterministic guarantees for data delivery has received considerable attention in the networking community [14, 15, 16]. A VBR stream is characterized by its worst-case demand on the network over any period of time. This characterization of the stream's behavior

*This work is supported in part by an NSF Career Award and by a grant from State of Texas Higher Education Board.

is used to calculate the worst-case delays possible through a switch when multiple streams pass through a switch. The service time for a packet through a switch is independent of the order in which the packets are served unlike at a disk where the order of service determines the seek time cost of the disk transfer. Similar worst case analysis at the disk is possible but likely to give less optimal results because of the higher seek time costs likely to be paid in any FCFS (First Come First Serve) approach.

Stream smoothing [17] has been proposed to reduce the variability of VBR streams. It has been shown that smoothing reduces the variance of demand and thus enables more number of streams to be supported at a network switch. Does smoothing help in disk retrieval as well? Previously proposed disk scheduling policies demonstrated the effectiveness of utilizing large blocks for data retrieval. Typically, these proposals have suggested retrieving data required for 0.25-1 seconds of display [11, 12]. Such a retrieval unit essentially smoothes data retrieval over the duration of the 'round'. Recently it has been shown that smoothing techniques can be applied across the rounds to achieve further reduction in the variability of stream's load on the disks [18].

With CBR streams, it has been shown that if the first block of the stream can be scheduled, then the entire stream can be provided guaranteed service [19]. To provide deterministic service to VBR streams, how does one guarantee that every block can actually be serviced? If such service can be provided, what is the best way to schedule the streams? If the stream is scheduled as soon as it can be scheduled, we can minimize the startup latency for that stream. However, this may hinder the scheduling of streams that may arrive in the future. This tradeoff between startup latency and the stream throughput of the system is investigated. We propose and evaluate a number of *stream scheduling* strategies for the same. This is different from a *disk scheduling* policy used for determining the order in which requests may be served at a disk.

It may be desirable to provide both deterministic service and statistical service to VBR streams in the same system. Deterministic service may be too expensive on the system's resources. A user may request for statistical service when a request for deterministic service may be denied due to lack of resources. How many more streams can be supported when statistical service is requested compared to that of deterministic service? We evaluate this issue as a function of the number of blocks that can be dropped.

This paper investigates a number of techniques for improving the throughput of VBR streams: (a) It proposes a technique for providing deterministic I/O service for VBR streams; (b) Considers the impact of stream startup latency on improving the stream throughput; (c) considers the impact of data smoothing on reducing the variability of data demands for improving the stream throughput; (d) considers the

impact of statistical guarantees on stream throughput and (e) suggests storing the peaks of popular streams in memory to improve stream throughput. Results from simulations quantify the impact of the above schemes.

Section 2 proposes a new approach for providing deterministic service guarantees at disk for VBR streams and also discusses how latency targets can be traded off for improving the throughput. Section 3 discusses how data layout, stream smoothing and statistical guarantees can improve the throughput. Section 4 presents a performance evaluation of these schemes based on trace-driven simulations. Section 5 summarizes our results and points out future directions.

2 Deterministic guarantees

Providing deterministic service for VBR streams is complicated by the following factors: (i) the load of a stream on the system varies from one round to the next, (ii) scheduling the first block doesn't guarantee that the following blocks of the stream can be scheduled. To ensure that all the blocks required by a stream can be retrieved, we can compute the peak rate of the stream and reserve enough disk bandwidth to satisfy the peak requirements of the stream. However, this approach results in considerable overestimate of required resources and results in supporting fewer streams. When multiple streams are served, can we exploit the statistical multiplexing that is possible while providing deterministic service?

It is assumed that the disk scheduling policy employs a round-based SCAN-like seek optimization [11, 12]. In such a policy, disk service is broken into fixed size time units called rounds or batches. Each round may span 0.25-1 seconds of time (based on earlier results in [11]). Requests are grouped into different rounds based on their deadlines. Requests within a round are served using SCAN or such similar seek optimizing policies. In this paper, we assume that the disk employs CSCAN, i.e., disk arm moves forward from the outermost request to innermost request during the round and jumps to the outermost request to start the next round.

In our approach, an application requiring service for a VBR stream supplies the I/O system with a trace of its I/O demand. This could be stored on the disk in the form of a special file or some attribute file. This data could be based on frame rate i.e., given on a frame to frame basis or could be more closely tied to the I/O system. Specifying the load on a frame basis is more flexible and the application doesn't have to be aware of how the I/O system is organized (block size or round size). If the I/O system's block size is known and the duration of each round is known, then the trace can be compacted by specifying the I/O load on a round by round basis in terms of the blocks. This requires less amount of data than the

frame by frame basis description. For example, a frame by frame trace may look like 83,888, 9,960, 10,008, 27,044, ... which indicates the number of bits of data needed to display each frame. If the round size is say 2 frames i.e., 1/12th of a second, and the I/O system uses a block size of 4KB, then the compacted trace would have $\lceil (83888 + 9960)/(4 * 1024 * 8) \rceil = 3$ in the first entry. The second entry would have $\lceil (10008 + 27044) - (3 * 1024 * 8 - 83888 - 9960)/(4 * 1024 * 8) \rceil = 1$ block. Hence, the equivalent compacted trace for the stream would be 3, 1, ... A 40,000 frame trace of the movie "Silence of the Lambs" (24 frames/second) requires 203,285 bytes on a frame by frame basis compared to a 3,333 byte description of the same movie when compacted with the knowledge of the round size of 0.5 seconds and a block size of 8KB. It is assumed that this information is available to the I/O system in either description and we will call this the *demand trace*.

The I/O system itself keeps track of the worst-case time committed for service in each round at each of its disks in the form of a *load trace*. Before a stream is admitted, its demand trace is combined with the load trace of the appropriate disks to see if the load on any one of the disks exceeds the capacity (committed time greater than the length of the round). The load trace of a system consists of load traces of all the disks over sufficient period of time. This requires the knowledge of the placement of blocks of the requesting stream. This information can be obtained from the file system.

A stream is admitted if its demand can be accommodated by the system. It is possible that the stream cannot be supported in the round the request arrives. The stream scheduling policy will look for a round in which this stream can be scheduled. Let $load[i][j]$ denote the load on disk i in round j . Let the demand of a stream be given by $demand[j]$ indicating the number of blocks to be retrieved by that stream in round j . Then, a stream can be admitted if there exists a k such that $load[i][j+k] + serv_time(demand[j]) \leq round\ time$, for all j , $i =$ disk storing data for round j , and k is the startup latency $\leq latency\ target$. In a CDL layout, where multiple disks may store the data required by a stream in a round, the above check needs to be appropriately modified to verify that each of the disks can support retrieval of a single block of data. The function $serv_time()$ estimates the worst-case service time required for retrieving a given number of blocks from a disk given the current load of the disk. This function utilizes the current load of the disk (number of requests and blocks) and the load of the arriving request to estimate the worst-case time required to serve the new request along with the already scheduled requests. A similar check can be applied against buffer resources when needed. An approach of keeping track of the load at the server (not at the individual disks as suggested here) is suggested for improving admission probabilities for VBR streams [20]. However, the approach in [20] cannot provide deterministic guarantees for VBR streams since all the

admitted streams may request data from the same disk in a single time slot (with CDL layout or CTL layout when all the streams are started at the same time). Our work here also explores other issues of data smoothing and statistical guarantees.

Given a latency target L and the length of the demand trace d , the admission controller requires at most Ld additions to determine if a stream can be admitted. In the worst case, for each starting round, the admission controller finds that the very last block of the stream cannot be scheduled. On an average, the admission controller requires less computation per stream. If necessary, latency targets can be reduced to limit the time taken by the admission controller.

This approach allows the system to exploit the statistical multiplexing that is possible when serving multiple VBR streams. Since the system ensures that the system has enough bandwidth to serve an arriving stream through its playback duration, deterministic guarantees can be provided. We will show through simulations that this results in significantly higher stream throughput than peak-rate allocation.

2.1 Stream scheduling

The number of streams that can be supported by the server is dependent on stream scheduling policies. An arriving request may not be scheduled immediately upon its arrival because of the already existing load on the system. By delaying the starting of the stream, we may be able to schedule a stream that may not otherwise be scheduled i.e., the startup latency can be traded off for stream throughput. How should streams be scheduled in a server? By greedily scheduling a stream as early as possible, we may impact the possibility of scheduling other streams in the future. If scheduling a stream immediately after its arrival saturates the capacity of a disk, that disk would be unavailable in that round for any other service and hence can affect schedulability of other streams. This issue is explored by evaluating a number of scheduling strategies.

All the scheduling algorithms discussed below use a *latency target* as a parameter. A stream is said to be unschedulable if it cannot be scheduled within a fixed time interval (specified by the latency target) after the arrival of the request. If a stream arrives at time t , all the slots within the time $(t, t + L)$ are considered for scheduling a stream, where L is the latency target. However, the order in which these slots are considered and the criterion for selection among the choices (if any) is determined by the stream scheduling policy. In the policies described below, irrespective of the starting point in the latency target window, all the available slots are considered for scheduling if necessary. It is noted that stream startup latency can be traded off to improve throughput in any system even if it doesn't employ the earlier proposed scheme for admission control. However, in this paper, we will restrict our study of this tradeoff to the context

of the earlier proposed admission control policy.

In *greedy scheduling*, a stream is scheduled as soon as it can be from the time the request arrives at the system. In *random start* policy, a stream is scheduled greedily from a random starting point within the latency target window. In *last scheduled* policy, a stream is scheduled greedily from the scheduled point of the last stream. In *fixed distance* policy, a stream is scheduled greedily from a fixed time away from the last scheduled stream's scheduled point. In *minimal load* policy, stream is scheduled at a point that minimizes the maximum load on any disk in the system. In *prime hopping* policy, instead of serially looking at the time slots from a starting point, slots a prime distance away are considered. For example, if the request arrives at time 0, a random starting point s is chosen. Then rounds, $s, s + p, s + 2p, s + 3p \dots$ are considered until the stream can be scheduled. Since p is prime, all the rounds within the latency target window will be considered.

Latency target impacts stream throughput in two ways. A larger target allows us to search more slots to find a suitable starting point for an arriving stream. A larger target also allows the peaks in the demand trace to be spread out more from each other and thus allowing a future stream to find enough I/O bandwidth to be scheduled.

Stream scheduling problem can be considered in two ways. In the first, given an existing load on the system, can an arriving stream be scheduled without disturbing the guarantees of already scheduled streams? This is the problem we consider in this paper. The scheduling decisions are made one at a time. Another interesting problem arises in capacity planning: can the system support the load of a given set of streams? This problem utilizes the information about all the streams at once to answer the question whether the system can support such a load (with required guarantees)? It can be shown that the stream scheduling problem is closely related to the bin packing problem which is known to be NP-hard [21].

3 Other approaches

In this section, we will discuss other schemes proposed earlier for improving the throughput of VBR streams.

3.1 Data layout

Data layout plays a significant role on the performance of disk access. It has been suggested by many researchers that video data should be striped across the disks for load balancing and to improve throughput available for a single data stream [5, 19, 6, 7]. The impact of data layout on the proposed scheme is investigated in this paper.

Data for a VBR stream can be stored in (i) Constant Data Length (CDL) units (ii) Constant Time Length (CTL) units. In CDL, data is distributed in some

fixed size units, say 64KB blocks. In CTL, data is distributed in some constant time units, say 0.5 seconds of display time.

With CDL layout, when data is striped across the disks in the system, data distribution is straightforward since each disk stores a block in turn. With CDL, data will be retrieved at varying rates based on the current rate of the stream. When data rate is high, data is requested more often. The variable rate of data retrieval makes it hard to combine such a policy with round-based seek optimizing policies. To make it possible to combine CDL layout with such seek optimizing policies, we consider data retrieval separately from the layout policy. Instead of retrieving one block at a time, display content for a constant unit of time is requested from the I/O system at once. If the data rate is high, this may result in requesting blocks from multiple disks in one time unit. Then data retrieval at these (possibly multiple) disks can employ round-based seek optimization. It is noted that the data required for display in a unit of time need not be a multiple of the I/O block size. Since the data retrieval is constrained by the I/O block size of the system, the needed data is rounded up to the next block.

In CTL layout, each disk stores the display content of the stream in a unit of time. In this paper, the data layout and retrieval is assumed to be constrained by the I/O block size. This is termed Block-constrained CTL data layout or BCTL in this paper. In BCTL, data distribution is harder since the amount of data stored on each disk depends on the data rate in that round and hence varies from disk to disk. However, data retrieval in a round-based scheme is easier since the data required by the stream in a round is located on a single disk. These issues have been addressed in [22, 10, 23] and the BCTL data layout is termed the hybrid scheme in [10].

3.2 Stream smoothing

The variability in resource requirements makes it difficult to schedule VBR streams. It has been suggested that buffering can be used to "smooth" a stream so as to reduce its burstiness or variability to increase its schedulability. A number of studies have shown its effectiveness for network scheduling, for example [17]. Disk scheduling also favors the use of larger blocks. When data is accessed in larger units of (time or data), the variability in accesses can be considerably reduced. Stream smoothing considers further prefetching of a stream data so as to reduce its burstiness. Fig. 1 shows an example of the resulting two traces of Lambs at a block size of 64KB and a round of 0.5 seconds. Smoothing reduces the peak demand of the stream and also reduces the variance of the stream's demand trace.

If larger request blocks are required by the efficiency of larger disk requests provide some form of smoothing, does additional smoothing provide benefits at a




Fig. 1. Effects of smoothing on a video stream.

server? To answer this question, we consider two separate streams for each movie, one unsmoothed and smoothed, and compare their performance. We used the optimal stream scheduling algorithm described in [17] for smoothing the movie streams. Smoothing reduces the resource demands by consider one stream at a time and the proposed admission control policy reduces resource demands by smoothing the resource requirement demands across multiple streams (through statistical multiplexing).

3.3 Statistical guarantees

Statistical guarantees are widely used when it is difficult to provide deterministic guarantees. With a 99% statistical guarantee, the server could afford to miss deadlines for about 1% of the blocks without violating guarantees. This laxity could allow more streams to be scheduled at the server than deterministic guarantees would allow. Since most video applications could tolerate the loss of a few blocks without an appreciable degradation in quality, statistical guarantees are widely used for video service.

How does the throughput improve when statistical guarantees are considered instead of deterministic guarantees? Since the proposed admission control allows statistical multiplexing of streams while providing deterministic guarantees, do statistical guarantees improve throughput over the proposed admission control policy? To answer this question, we consider statistical guarantees up to 95%.

4 Performance Evaluation

4.1 Simulations

We evaluated a number of the above issues through trace-driven simulations. A system with 8 disks is simulated. Each disk is assumed to have the characteristics of a Seagate Barracuda drive [24]. The disk drive characteristics are shown in Table 1. Each disk in the system maintains a load table that depicts the load of that disk into the future. An arriving stream presents its demand trace to the simulator. Based on the data layout strategy employed, the stream's demand trace is combined with the load trace of the system. The starting point of the stream and the order in which candidate slots are tried for admitting a stream is determined by the stream scheduling policy. CDL and BCTL data layout strategies are considered. A number of stream scheduling policies are explored. Data block size on the disk is varied from 32KB to 256KB. In BCTL, it is assumed that each stream pays a latency penalty at a disk. In CDL, a stream pays at most one latency penalty at each disk per round. For example, if the stream requires 1 block each from disks 1, 2 and 3, then that stream pays a latency penalty at disks 1, 2 and 3 in that round. If the stream requires 10 blocks in a round from the 8 disks in the system, it pays a latency penalty at each disk. This is based on the assumption that the blocks retrieved in a round for a stream are stored contiguously on the disk.

Stream scheduling policies do not alter already scheduled streams to admit a new stream. Also, when requests arrive at the same time, admission decisions are made one stream at a time. All the stream requests are assumed to arrive in time slot 0. If the stream requests are assumed to arrive randomly over time,

Table 1. Disk characteristics.

Parameter	Value
Zero Seek time	0.85 ms
Avg. Seek time	9.0 ms
Max. Seek time	18.0 ms
Min. Transfer rate	11.5 MB/s
Max. Transfer rate	17.5 MB/s
Full Rev. time	8.33 ms
Num. cylinders	6311

Table 2. Characteristics of traces.

Stream Name	Mean KB/sec	Sta. Dev.
Lambs	171.32	58.33
Term	255.54	50.92
News	484.31	108.86
Asterix	523.79	124.50

more streams can be admitted. However, to study the worst-case scenario, we assumed that all the requests arrive at once. The simulator tries to schedule as many streams as possible until a stream cannot be scheduled. The number of streams scheduled is the stream throughput. This process is repeated for all the scheduling policies and block sizes. Four different video streams are considered in our study as explained below.

4.1.1 Traces

MPEG traces from University of Wuerzburg [25] were used in this study. From the frame by frame trace of the movie, we constructed several versions of the demand trace for each movie. For this study, we used four separate MPEG traces. These traces are named *Lambs* (for a segment of the movie "Silence of the lambs"), *Term* (for a movie segment of the movie "Terminator"), *News* (for a news segment trace), *Asterix* (for a segment of Asterix cartoon). Each trace contained 40,000 samples at a frame rate of 24 frames per second. A mixed workload based on these traces is also constructed. During simulations, with equal probability, one out of the four traces is selected for scheduling i.e., the workload consisted of a random mix of these four traces with each tracing being selected with an equal probability. Each trace has a different bit rate and different mean and variance characteristics and these are shown in Table 2. The traces obtained were of fairly low bit rates. To simulate better quality streams, we assumed that the bit trace is actually a byte trace without changing the relative weights of the frames.

A block size of 32 KB, 64KB, 128KB or 256 KB and 0.5 seconds of round time are used to convert the frame trace into a compact demand trace for each movie segment. With the choice of four block sizes,

we get four different compact demand traces for each stream. These four different traces are used to study the impact of the block size on the results. Next we smoothed the trace of each stream by assuming that a memory buffer of 8 blocks is available and then converted the smoothed trace into a compact demand trace in terms of blocks. For each block size, we have an unsmoothed demand trace and a smoothed demand trace for a total of 8 separate traces for the 4 block sizes considered in this study. In simulations, it is assumed that the first block of each data stream is stored on a different disk. Data is striped across the eight disks in a round-robin order based on either CDL or BCTL data layout policy.

4.2 Results

Fig. 2 shows the impact of the proposed admission control policy on the stream throughput of the four different data streams separately. The results shown in Fig. 2. are obtained with random-start stream scheduling policy. For both the data layouts considered, the proposed admission control policy performs significantly better than a policy based on peak rates of the streams. The proposed approach achieves 130%-195% more stream throughput than a method based on peak rate allocation. The higher throughput is achieved by exploiting the statistical multiplexing across multiple streams.

For the block sizes considered, the BCTL data layout strategy performs better than the CDL strategy. Since CDL results in paying a rotational latency penalty for each block compared to each time slice in BCTL, CDL performs worse than BCTL. As the block size is increased a stream fetches less number of blocks in a round and hence CDL tends to be more efficient at larger block sizes. The stream throughput for CDL improves significantly for all the data streams as the block size is increased from 32 KB to 256 KB. At smaller block sizes, the difference in performance between BCTL and CDL is larger and as the block size is increased, the difference decreases. At 256KB block size, both schemes achieve nearly the same stream throughput for Lambs and Term workloads. For these two data streams, at this block size, data retrieval requires more than a single block in very few rounds and hence the performance is nearly the same whether CDL or BCTL strategy is adopted. However, in Asterix and News, data rates are high and even at a block size of 256KB, considerable number of rounds require multiple block retrievals. As a result, we find significant difference in performance even at a block size of 256KB. The stream throughput drops slowly for BCTL as the block size is increased. This is due to effects of larger quantization of service allocation for a request.

Fig. 3 shows the improvements possible through stream smoothing. To avoid data clutter, we only show the results with the Lambs workload, but similar trends were observed across all the four video




Fig. 2. Performance of proposed admission control policy.

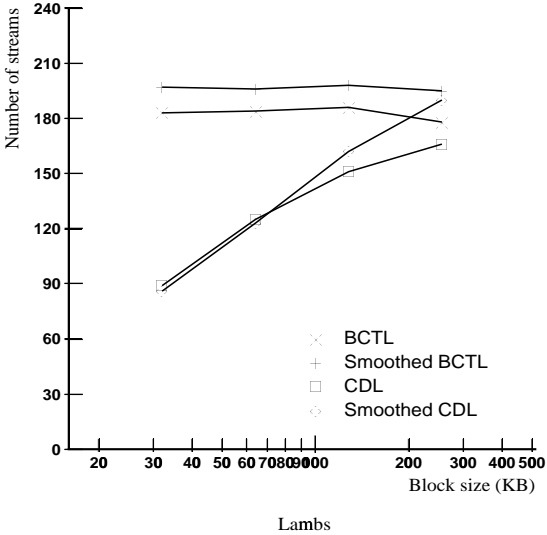


Fig. 3. Impact of smoothing.

streams. Smoothing improves performance with both data layout strategies. This is as expected since smoothing reduces the variability of demand which in turn helps in scheduling more streams at the I/O system. Smoothing can improve stream throughput by up to 17% for these workloads on top of the proposed admission control policy. Since we assumed that 8 blocks of buffers were available for smoothing, we see better benefits due to smoothing at larger block size.

From figures 2 and 3, BCTL is seen to offer better performance over CDL approach. However, at larger blocks, the performance differences are less significant. CDL approach is more flexible since it doesn't tie data distribution closely with the application characteristics (data consumption rate in each frame). To obtain good performance out of CDL, the block size should be large enough such that multiple blocks of data are not requested by a stream in most rounds. If the system supports streams at different data rates, then the block size and round size should be chosen such that the stream with highest data rate would retrieve one block in most rounds. This suggests employing larger block sizes than previously suggested by the disk scheduling policies [11].

Fig. 4. shows the disk utilization by the video streams as a function of time. The figure shows the load at one of the eight disks in the system with a mixed workload and *minimal load* stream scheduling policy. Even though the average utilization is 66%, for several seconds, the disk is nearly 100% busy between rounds 1800 and 2600. If we allowed the video streams to occasionally utilize the full 100% I/O bandwidth of the system while maintaining the average utilization below say 65%, the aperiodic requests could get starved for service for long periods of time (in this case for 400 seconds). Hence, this is unacceptable




Fig. 4. Disk utilization by VBR streams.

in a system that has to support both types of requests. This result shows that the system has to reserve certain minimum amount of I/O bandwidth for aperiodic requests to provide reasonable response to those requests.

Fig. 5. shows the impact of stream scheduling policies on the stream throughput at various block sizes. The results in Fig. 5. are for a mixed workload and a latency target of 300 rounds. Greedy policy achieves the least stream throughput in both data layout schemes. It is observed that the minimal load policy achieves high stream throughput consistently in BCTL data layout. However, minimal load policy doesn't perform as well with CDL data layout. Prime hopping, fixed distance and random start achieve nearly the same stream throughput. Minimal load policy achieves on an average 15% better stream throughput than these three policies with BCTL layout. It is observed that stream scheduling policy has a significant impact on performance. For example, minimal load policy improves performance by about 80% compared to greedy policy at a block size of 32KB in BCTL data layout. This shows the importance of studying the stream scheduling policies.

Fig. 6. shows the startup latencies achieved by different stream scheduling policies. Greedy policy, by its nature, achieves the smallest startup latency. However, as observed earlier, it also results in lower stream throughput. The policies based on randomness, random start, prime hopping and fixed distance achieve average startup latencies close to 150, which is half of the latency target of 300 rounds considered for these results. Minimal load and Last scheduled achieve better latencies than these policies for both the data layouts. Based on achieved stream throughput and startup latencies, minimal load and last scheduled are better candidates than the other policies. Last

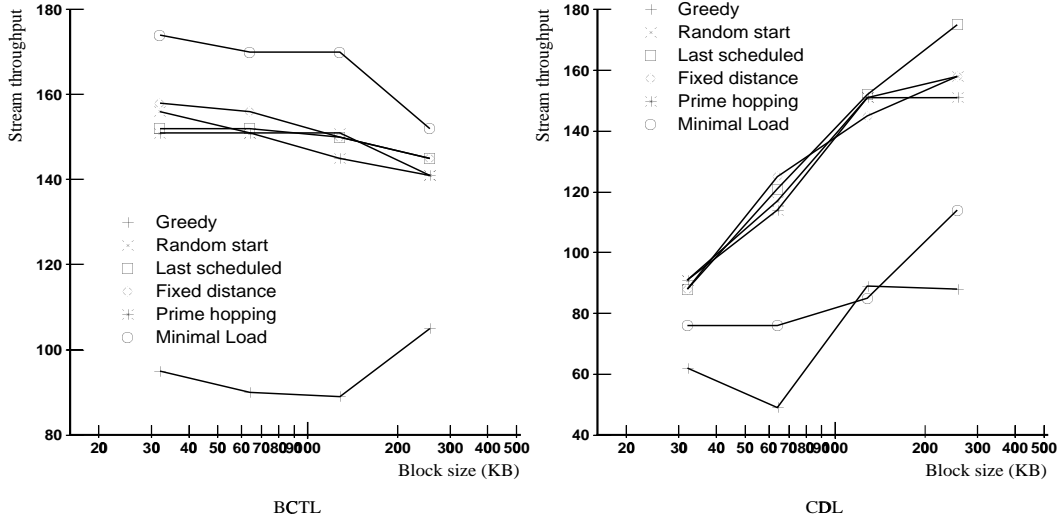


Fig. 5. Impact of stream scheduling.

scheduled has a more consistent performance across the two data layout strategies. (*Note to reviewers: Minimal load seems to be getting stuck close to the starting point with CDL data layout and we are investigating ways to improve this policy for CDL data layout.*)

Fig. 7. shows the impact of statistical guarantees and latency targets on stream throughput. Instead of requiring that every block of data be retrieved in time, we allowed a fraction of the blocks for each stream to miss deadlines or not be provided service. This fraction is varied among 0.1%, 0.2%, 0.5%, 1.0%, 2.0% and 5.0% at various latency targets. It is observed as more blocks are allowed to be dropped, it is possible to achieve more throughput compared to deterministic guarantees. Stream throughput can be improved by up to 20% by allowing 5% of the blocks to be dropped. Dropping blocks is more effective at lower latency targets than at higher latency targets. For example, dropping up to 2% of the blocks improves the stream throughput by 14.5% at a latency target of 100 rounds compared to an improvement of 6% at a latency target of 1000 rounds. It is also observed that stream throughput can be increased by relaxing the latency targets.

Fig. 7. also shows the tradeoffs possible between latency targets and the number of blocks allowed to be dropped. At a latency target of 100 rounds, 152 streams can be provided deterministic service. To achieve higher stream throughput, we can either increase the latency target or allow blocks to be dropped. For example, when we increase the latency target to 1000 rounds, 179 streams could be scheduled without dropping any blocks. However, to achieve the same throughput at a latency target of 100 rounds, more than 2% of the blocks have to be dropped. Hence, desired throughput can be achieved either

by allowing larger latency targets or by allowing a fraction of the blocks to be denied service.

Fig. 7. also shows the impact on the average disk utilizations as a function of the statistical allowances and latency targets. As higher statistical allowances are made, the disk utilizations are improved as more number of streams are supported. As latency targets are increased from 100 rounds, average disk utilizations first increase and then decrease to lower levels. As latency targets are increased, an arriving stream finds more choices to find a suitable starting spot to utilize the available bandwidth. However, as the latency targets are increased further, the streams are scheduled farther and farther into the future and hence result in decreasing disk utilizations.

Usually considered statistical guarantees of 99% (i.e., dropping 1% of blocks) did not provide significant improvements in stream throughput compared to deterministic guarantees. In our measurements, the improvements were less than 6% for all the latency targets except 100 which achieved an improvement of 11%. We observed that most of the dropped blocks tended to be dropped in a small range of time. Even though a 30 minute movie drops only a few blocks during its playback, the dropped blocks tended to be clustered over a 2 minute window instead of over the entire movie. This observation that the dropped blocks tended to be clustered together over shorter time windows leads us to the following idea. It may be possible to drastically improve the stream throughput of a popular movie if the blocks during its peak demand slots are stored in memory. By storing only a few blocks instead of the entire movie, we require less memory. If the movie is popular, several streams will get benefit out of storing these blocks in memory. These ideas will be explored in the future work.

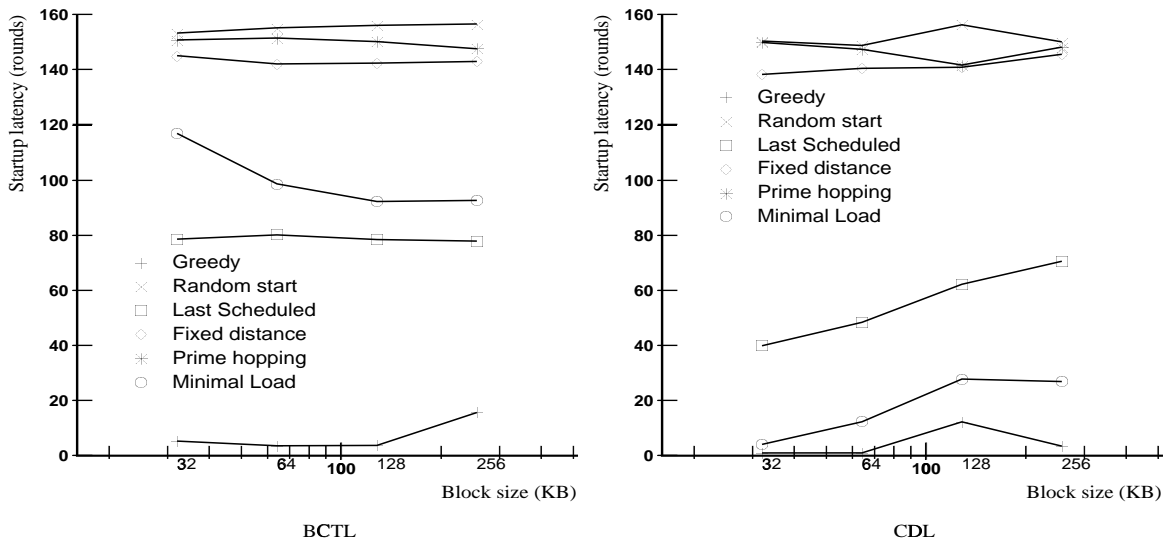


Fig. 6. Average startup latency.

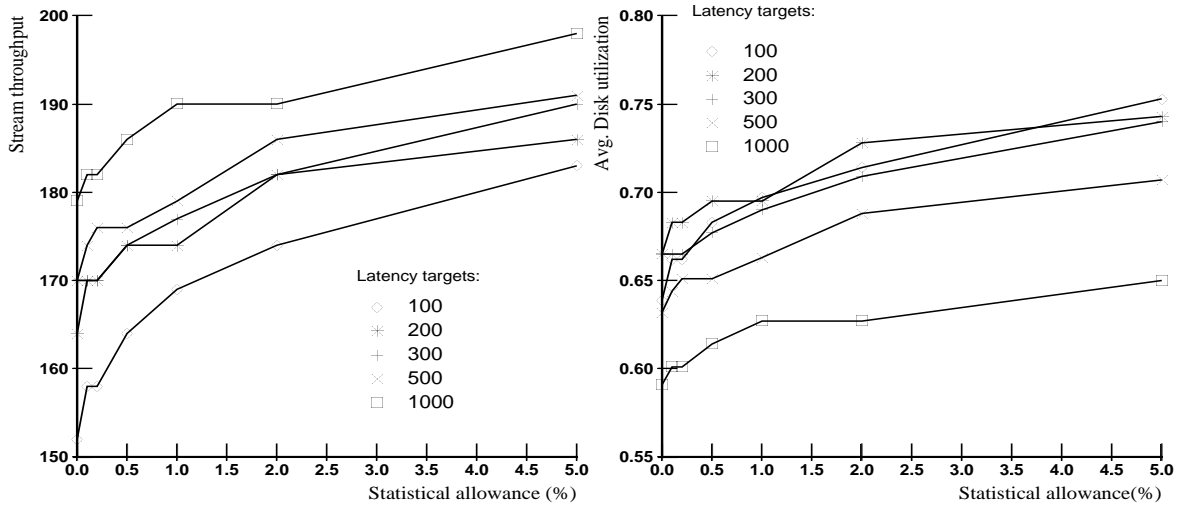


Fig. 7. Effect of statistical allowances.




Fig. 8. Impact of arrival rates.

Fig. 8. shows the impact of arrival rates over number of streams scheduled and the average disk utilization. As expected, as the inter arrival time increases, more streams can be scheduled. However, the average disk utilization decreases with increased inter-arrival times. Increasing statistical allowances improve the utilized disk bandwidth.

5 Conclusions

In this paper, We investigated a number of techniques for improving the throughput of VBR streams. We presented an approach for providing deterministic disk service for VBR streams. The proposed approach exploits statistical multiplexing across multiple streams while providing retrieval guarantees. We have shown through simulations that the proposed approach improves stream throughput by 130%-195% over peak-rate allocation. Several stream scheduling strategies were presented and evaluated. Several of these proposed policies are shown to improve performance significantly compared to the greedy policy. We also evaluated the performance impact of a number of issues such as data layout, stream scheduling, startup latency and statistical guarantees. We showed that startup latency is an effective tradeoff parameter for improving stream throughput. Stream smoothing is shown to improve stream throughput marginally on top of the proposed admission control policy. For the workloads considered, dropping 1% of the blocks in addition to the proposed admission control policy, did not provide significant improvement in stream throughput. Our results also showed the need for reserving certain minimum I/O bandwidth for aperiodic

requests.

References

- [1] H. M. Vin and P. V. Rangan. Designing file systems for digital video and audio. *Proc. of 13th ACM Symp. on Oper. Sys. Principles*, 1991.
- [2] D. Anderson, Y. Osawa, and R. Govindan. A file system for continuous media. *ACM Trans. on Comp. Systems*, pages 311–337, Nov. 1992.
- [3] C. Martin, P. Narayanan, B. Ozden, R. Rastogi, and A. Silberschatz. The Fellini multimedia storage server. in *Multimedia Information Storage and Management*, Ed: S.Chung, KluwerPublishers, 1996.
- [4] D. J. Gemell and S. Christodoulakis. Principles of delay-sensitive multimedia storage and retrieval. *ACM Trans. on Info. Systems*, pages 51–90, 1992.
- [5] R. Haskin. The shark continuous-media file server. *Proc. of IEEE COMPCON*, Feb. 1993.
- [6] Microsoft. The tiger video server. *Microsoft Press Release*, Apr. 1994.
- [7] A. Laursen, J. Olkin, and M. Porter. Oracle media server: providing consumer based interactive access to multimedia data. *Proc. of SIGMOD*, pages 470–477, 1994.
- [8] T. Niranjan, T. Chiueh, and G. A. Schloss. Implementation and evaluation of a multimedia file system. *Proc. of IEEE Conf. on Multimedia*

- Computing and Systems*, pages 269–276, June 1996.
- [9] H. Vin, P. Goyal, and A. Goyal. A statistical admission control algorithm for multimedia servers. *Proc. of ACM Multimedia*, Nov. 1994.
- [10] E. Chang and A. Zakhor. Scalable video data placement on parallel disk arrays. *Proc. of SPIE Symp. on Elec. Imaging Sci. and Tech.*, Feb. 1994.
- [11] A. L. Narasimha Reddy and Jim Wyllie. Disk scheduling in a multimedia I/O system. *Proc. of ACM Multimedia Conf.*, Aug. 1993.
- [12] P. S. Yu, M. S. Chen, and D. D. Kandlur. Grouped sweeping scheduling for dasd-based multimedia storage management. *Multimedia Systems*, 1:99–109, 1993.
- [13] R. T. Ng and R. Dilek. Statistical modeling and buffer allocation for MPEG streams. in *Multimedia Information Storage and Management*, Ed: S. M. Chung, Kluwer Academic Publishers, 1996.
- [14] D.E. Wrege, E. W. Knightly, H. Zhang, and J. Liebeherr. Deterministic delay bounds for VBR video in packet-switching networks: fundamental limits and practical trade-offs. *IEEE/ACM Trans. on Networking*, pages 352–362, June 1996.
- [15] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated service networks: The multiple node case. *Proc. IEEE INFOCOMM'93*, pages 521–530, Mar. 1993.
- [16] R. Cruz. A Calculus for Network Delay, Part I: Network elements in isolation. *IEEE Trans. on Information Theory*, vol.37, no.1, pages 114–131, Jan. 1991.
- [17] J. Salehi, Z. L. Zhang, J. Kurose, and D. Towsley. Supporting stored video: reducing rate variability and end-to-end resource requirements through optimal smoothing. *Proc. of ACM SIGMETRICS*, May 1996.
- [18] S. Sahu, Z. L. Zhang, J. Kurose, and D. Towsley. On the efficient retrieval of VBR video in a multimedia server. *Proc. of IEEE Int. Conf. on Multimedia Computing and Systems*, Jun. 1997.
- [19] A. L. N. Reddy. Scheduling and data distribution in a multiprocessor video server. *Proc. of Int. Conf. on Multimedia Computing and Systems*, pages 256–263, 1995.
- [20] E. Chang and A. Zakhor. Cost analyses for vbr video servers. *Proc. of SPIE Multimedia Computing and Networking*, pages 381–397, Jan. 1996.
- [21] M. R. Garey and D. S. Johnson. Computers and intractability: A guide to the theory of np-completeness. *W.H.Freeman & Co., New York, NY*, 1979.
- [22] H. Vin, S. S. Rao, and P. Goyal. Optimizing the placement of multimedia objects on disk arrays. *Proc. of IEEE Conf. on Multimedia Computing and Systems*, pages 158–165, May 1995.
- [23] M. Buddhikot and G. Parulkar. Efficient data layout, scheduling and playout control in MARS. *ACM Journal on Multimedia Systems*, pages 199–211, 1997.
- [24] Seagate Corp. Disk drive product info. <http://www.seagate.com/>, 1997.
- [25] O. Rose. Mpeg trace data sets. <ftp://info3.informatik.uni-wuerzburg.de>, 1995.