Maximizing Buffer and Disk Utilizations for News On-Demand

Raymond T. Ng and Jinhai Yang
(rng@cs.ubc.ca) (jyang@cs.ubc.ca)
Department of Computer Science
University of British Columbia
Vancouver, B.C., V6T 1Z4, Canada

Abstract

In this paper, we study the problem of how to maximize the throughput of a multimedia system, given a fixed amount of buffer space and disk bandwidth both pre-determined at design-time. Our approach is to maximize the utilizations of disk and buffers. We propose doing so in two ways. First, we analyze a scheme that allows multiple streams to share buffers. Our analysis and preliminary simulation results indicate that buffer sharing could lead to as much as 50% reduction in total buffer requirements. Second, we develop two prefetching strategies: SP and IP. As will be demonstrated by SP, straightforward prefetching is not effective at all. In contrast, IP, which prefetches more intelligently than does SP, could be valuable in maximizing the effective use of buffers and disk. Our preliminary simulation results show that IP could lead to a 40% improvement in throughput.

1 Introduction

With the advances in networking, storage, and I/O interface technologies, providing effective multimedia support in database management systems has become a topic of great interest and value. To support audio and video data, multimedia database management systems need to deal with several tough issues. First, audio and video data are delay-sensitive. As recording and playback of video and audio data are continuous operations, a management system, once starts displaying audio or video data, must guarantee that enough resources are allocated so that the *continuity* and real time requirements are not violated. Second, (even compressed) audio and video data consume large amounts of system resources – primarily storage space and bandwidth. Third, a multimedia object may consist of multiple components: audio, video and text. It is the responsibility of the management system to ensure that these multiple streams can be synchronized during retrieval.

Many excellent studies regarding the storage and retrieval of audio and video data have been conducted, such as those reported in [1, 2, 3, 4, 9, 10, 12, 13]. With respect to the topic area of this paper, these studies can be grouped into two major categories. The first group is primarily concerned with intelligent disk scheduling. Studies in this group include the sweeping scheme proposed by Chen, Kandlur and Yu [2], the sorting-set algorithm developed by Gemmel [3], the SCAN-EDF strategy designed by Reddy and Wyllie [10], and the hard real-time approach analyzed by Tindell and Burns [12]. The second group deals with constrained block allocation, which limits the distance between successive blocks of a multimedia stream. Studies in this group include the scattering parameter approach developed by Rangan and Vin [9], the cluster strategy introduced by Gemmel and Christodoulakis [3, 4], and the audio data placement work of Yu et. al. [13]. To a very large extent, most of these proposals aim to minimize seek latencies so as to satisfy the continuity requirements of multimedia streams. And most of them are developed from

Proceedings of the 20th VLDB Conference
Santiago, Chile, 1994

a design perspective.

Quite complementary to the problems addressed in the studies mentioned above, the problem we consider here is concerned with the dynamic aspect of a multimedia system. More specifically, given a fixed amount of buffer space and disk bandwidth, both predetermined at design time, we study how to maximize the throughput of a multimedia system, and minimize the response time of queries (with the guarantee that all continuity requirements will be satisfied). For a system with fixed disk bandwidth and buffer space, the response time of queries are primarily governed by the utilization of disk and buffers. Thus, our approach is to utilize buffers and the disk as effectively as possible. In particular, in this paper, we will report:

- a scheme that allows multiple streams to share buffers. We will give an analysis on the benefit of buffer sharing, which could lead to a 50% reduction in total buffer requirements. We will also present preliminary simulation results providing further evidence showing the effectiveness of buffer sharing.
- two prefetching strategies: SP and IP. As will be shown by Strategy SP, straightforward prefetching may not be effective at all. In contrast, Strategy IP, which prefetches more intelligently than SP, could be very valuable in maximizing disk and buffer utilizations, as well as system throughput. Our preliminary simulation results indicate that IP could lead to a 40% improvement in throughput.

On first sight, the class of multimedia systems most amenable to the techniques proposed here is the class of news on-demand systems (e.g. [6]). This is because such systems normally have non-uniform, asynchronous arrivals of queries, and the lengths of queries are usually not long (e.g. $\leq$ 5 minutes). However, in Section 5.4, we will argue that the proposed techniques can also be applicable to multimedia systems other than news on-demand.

The organization of the paper is as follows. Section 2 presents a preliminary analysis on periodic retrieval of multiple streams, and gives several basic equations needed in later analyses. Section 3 presents an analysis on buffer sharing. Section 4 introduces and analyzes the two prefetching strategies: SP and IP. Section 5 presents preliminary simulation results, followed by discussions and conclusions.

2 Preliminary Analysis: Periodic Retrieval of Multiple Streams

As observed in [4, 9], for various performance reasons, the most efficient way to process multiple streams simultaneously is to interleave the reading of the streams in a cyclic fashion. In this section, we provide a preliminary analysis of this situation.

Symbol	Meaning of Symbol
B_{max}	maximum number of buffers available
B	total buffers needed by n streams
B_{shar}	total buffers needed with buffer sharing
B_i	buffer consumption of Stream S_i
Bl	block size in non-contiguous placement
P	total consumption rate of n streams
P_i	consumption rate of Stream S_i
P_i^{pft}	consumption rate of S_i after prefetching
R	maximum disk reading rate
S_i	the i-th stream
s	total switching time within a cycle
$s_{i,j}$	switching time between Streams S_i and S_j
t	length of a cycle
t_i	reading time for S_i within a cycle
T_i	length of Stream S_i (in seconds)
ρ	disk utilization

Figure 1: Meanings of Symbols Used

Let there be n multiple streams denoted by $S_1, \ldots, S_n$. Let the consumption rate [1] of Stream S_i be P_i, and the amount of time reading S_i in each period be t_i. Then if $s_{i,j}$ denotes the seek (or switching) time from S_i to S_j, we have: $t_1 + \ldots + t_n + s_{1,2} + \ldots + s_{n,1} \leq t$, where t denotes the total length of the cycle. To simplify notations, let $s = s_{1,2} + \ldots + s_{n,1}$. Then the disk utilization, ρ, is given by:

$$\rho = \frac{t_1 + \ldots + t_n + s}{t} \tag{1}$$

Figure 1 summarizes the meanings of the symbols to be used in this paper.

Now let us take a closer look at each Stream S_i. The analysis below assumes that apart from the seek required for switching from S_{i-1} to S_i, no extra seek is needed throughout time t_i when S_i is being read. This can be achieved by using the technique of storing data in clusters proposed in [3], or by storing data contiguously (e.g. such as in a spiral optical disk). In [8], we discuss how to relax this assumption to handle other situations of data placement.

Within each period, the total amount of data consumed by S_i is $t * P_i$, and the amount read for S_i is $t_i * R$, where R is the maximum disk reading rate. Thus, the continuity requirement of S_i can be expressed as:

$$t_i * R \geq t * P_i \tag{2}$$

[1] The consumption rate refers to the rate the data obtained from disk are consumed. For an uncompressed stream, its consumption rate is the same as its playback rate.

However, in order to reduce the number of buffers used for each stream, we have:

$$t_i * R \;=\; t * P_i \tag{3}$$

From Equation 3, it is easy to see that $\frac{t_i}{t_j} = \frac{P_i}{P_j}$. In other words, to minimize buffer consumption, the reading time for each stream should be proportional to its consumption rate. Let P denote the total consumption rate, i.e. $P = P_1 + \ldots + P_n$. Then by combining Equations 1 and 3, t_i can be determined by:

$$t_i \;=\; (t * \rho - s) * \frac{P_i}{P} \tag{4}$$

The above equation gives the amount of reading time for S_i in terms of t, the length of the cycle. In the following, we establish a lower bound on t, by combining Equations 2 and 4:

$$t \;\geq\; \frac{s * R}{R * \rho - P} \tag{5}$$

This equation leads to two interesting observations. First, the equation is valid only if $(R * \rho - P) > 0$. Even if the disk utilization ρ is set to the maximum 1, it is necessary that $R > P$. This is the most obvious admission control criterion. That is, without violating their continuity requirements, a system cannot admit so many streams that their total consumption rate P exceeds the disk bandwidth. In Section 4, we will show how this constraint can be relaxed by prefetching.

Second, t is inversely proportional to ρ. In other words, the longer the length of the period, the less utilized the disk becomes (for the n streams). This is because as t increases, the proportion of time wasted in switching (i.e. $\frac{s}{t}$) within every cycle becomes smaller. In other words, a longer period corresponds to a higher percentage of useful work (i.e. data transfer) done by the disk, and the disk becomes more effective. Hence, the proportion of the time when the disk is idle becomes higher. In Section 4, we will show how to make use of this relationship between t and ρ to maximize prefetching.

3 Analysis on Buffer Consumption

Thus far, we have analyzed the handling of multiple streams primarily from the viewpoint of disk bandwidth allocation. There is, however, another dimension: the allocation of buffers. In this section, we will first give equations specifying the buffer requirements, based on the analysis presented in the previous section. Then we will analyze how *sharing* of buffers among streams can minimize total buffer consumption and maximize buffer utilization.

3.1 Buffer Requirements of Multiple Streams

Recall from the above analysis that the basic strategy to support multiple streams simultaneously is that for each Stream S_i, enough data of S_i must be read in time t_i to cover the consumption of S_i for time t. To achieve this, buffers are needed for S_i. In particular, the maximum number of buffers is needed right after S_i has just finished reading. Thus, the number of buffers required by S_i is: $B_i = t_i * R - t_i * P_i$. By substituting Equation 4 into the above, we get:

$$B_i \;=\; P_i * (R - P_i) * \frac{t * \rho - s}{P} \tag{6}$$

Thus, the total buffer requirements for the n streams is:

$$B = \sum_{i=1}^{n} B_i \;=\; \frac{t * \rho - s}{P} * \sum_{i=1}^{n} P_i * (R - P_i) \tag{7}$$

Two observations can be drawn from the above equation. First, it is obvious from the equation that the longer the period length t, the higher the value of B is. Second, if B_{max} is the maximum number of buffers available in the system, it is necessary that $B \leq B_{max}$. By substituting Equation 7 into $B \leq B_{max}$, we get an upper bound of the cycle length t:

$$t \;\leq\; \frac{B_{max} * P}{\rho * \sum_{i=1}^{n} P_i * (R - P_i)} + \frac{s}{\rho} \tag{8}$$

This equation can be combined with Equation 5 to provide the following admission control policy.

Admission Control Let $S_1, \ldots, S_{n-1}$ be all the streams in the current cycle, and S_n be the stream to be decided whether admission is possible.

1. Compute the lower bound (of t) based on Equation 5 and the upper bound based on Equation 8.
2. If the lower bound is strictly greater than the upper bound, then it is not possible to add S_n without violating continuity requirements.
3. Otherwise, S_n can be admitted to form a new cycle, and any value between the lower and upper bound can be chosen as the length of the new cycle. □

In Section 4, we will return to this issue of picking a value for t, and analyze in greater details how to do that to maximize prefetching.

3.2 Buffer Sharing and its Benefit

As defined in Equation 7, the total buffer requirement of n streams is based on the assumption that each

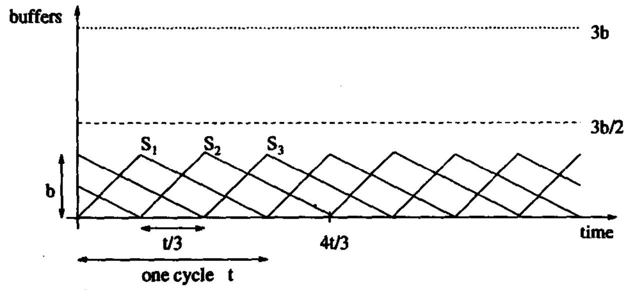

Figure 2: Buffer Sharing for 3 Streams with Identical Consumption Rates

stream S_i occupies B_i buffers within each cycle. However, as shown in Figure 2, S_i does not need all B_i buffers at all times. In fact, S_i's buffer requirement can be less than B_i, for example when $S_{i+1}, \ldots, S_n$ require their maximum number of buffers. Thus, a simple way to minimize total buffer consumption and thus to maximize buffer utilization is to allow the n streams to share buffers.

Figure 2 shows a simple situation when there are 3 streams S_1, S_2, S_3 in the cycle, all of which has the same consumption rate. Thus, by Equation 4, each stream has an equal amount of reading time, i.e. same t_i. Since the cycle length t is normally much larger than the total switching time s, Figure 2 shows the simplified situation when $t_i = t/3$. Let us consider the total buffer requirement at time $4t/3$, at which point S_1 has just finished reading and requires b buffers, the maximum number of buffers that it ever needs. S_2, which is about to start reading, has run out of data. Thus, the buffer requirement of S_2 is 0. As for S_3, there were b buffers at time t, but at time $5t/3$, all the data in those buffers will be consumed. Thus, at the current time $4t/3$, S_3 needs $b/2$ buffers. Hence, the total number of buffers required by all 3 streams is $b + 0 + b/2 = 3b/2$. Note that if all the streams have identical consumption rates, their total buffer requirement does not change with time. Thus, $3b/2$ buffers are all the 3 streams need. However, without buffer sharing, $3b$ buffers are required. Thus, buffer sharing gives a 50% reduction in total buffer consumption.

An analysis on buffer sharing for the general case when there are n streams with heterogeneous consumption rates $P_1, \ldots, P_n$ involves finding the time point within a period when the total buffer requirement reaches the maximum. This is necessary because this maximum is no longer constant when $P_1, \ldots, P_n$ are not all the same. Because such an analysis is very lengthy and complicated, we consider it beyond the scope of this paper. In the following, we will only analyze the situation when there are n streams with identical consumption rates. Since the consumption rates are the same, the individual buffer requirement B_i is the same, which is equal to b say. Similarly, the reading time t_i for each stream is the same, say equal to t_0. Now let us consider the time when S_n has just finished reading. The following table shows the buffer requirement of each stream at that point.

Streams	S_1	S_2	S_3	$\ldots$	S_n
Buffers needed	0	$\frac{1}{n-1}b$	$\frac{2}{n-1}b$	$\ldots$	$\frac{n-1}{n-1}b$

First, S_n has just finished reading, thus requiring all b buffers. S_1 is about to start reading. Thus, it has 0 buffers of data at this point. S_2, at an earlier point in time, had b buffers of data which are supposed to cover the consumption of S_2 for a period of $(n-1)*t_0$. At the point when S_n has just finished reading, $(n-2)*t_0$ has elapsed, or alternatively, S_2 will run of data t_0 seconds later. Thus, the current level of buffered data for S_2 is $\frac{t_0}{(n-1)*t_0}b = \frac{1}{n-1}b$. Similarly, it is not difficult to see that the current level of buffered data for S_3 is $\frac{2}{n-1}b$. Hence, the total number of buffers needed is:

$$B_{shar} = \sum_{i=1}^{n} \frac{i-1}{n-1}b = \frac{n}{2}b \tag{9}$$

In this case, without buffer sharing, the total number of buffers required is $B = nb$. Thus, buffer sharing reduces total buffer consumption by 50%.

Example 1 Consider a homogeneous set of streams whose consumption rate is 240KB per second. (This is based on 24 frames per second where each frame is JPEG compressed to 10KB [11].) Given a disk whose maximum reading rate is 1000KB per second, 4 streams can be supported simultaneously, provided that there are enough buffers. Let the total switching time be $s = 0.1$ secs. Furthermore, let us pick the minimum cycle length, which corresponds to $\rho = 1$. Then by Equation 5, $t = 2.5$ secs. By Equation 6, the maximum buffer requirement for each stream is $b = 456$KB. Thus, without buffer sharing, about 2MB of buffer space is needed. But with buffer sharing, only 1MB is needed. Alternatively, if the system only has 1MB of buffer space, the number of streams that can be supported simultaneously without buffer sharing is only 2. With buffer sharing, the system can double the throughput and support all 4 streams. □

The above analysis assumes that the disk utilization ρ is equal to 1. To take disk utilization into account, we generalize the above table that shows the buffer requirement of each stream at the point after S_n has finished reading to become:

Streams	S_1	S_2	$\ldots$	S_n
Buffers	cb	$(c + \frac{\rho}{n-\rho})b$	$\ldots$	$(c + \frac{(n-1)\rho}{n-\rho})b$

where $c = \frac{1-\rho}{1-\rho/n}$. A simple summation yields: $B_{shar} = \frac{2n-n\rho-\rho}{2(n-\rho)} * nb$. Now according to Equation 6, b is equal to $(R - \frac{P}{n}) * \frac{t*\rho-s}{n}$. Thus, the full equation is:

$$B_{shar} = (R - \frac{P}{n}) * (t * \rho - s) * \frac{2n - n\rho - \rho}{2(n - \rho)} \quad (10)$$

This equation can replace Equation 7 (and thus Equation 8) in the admission control test shown in Section 3.1. In Section 5, we will present preliminary simulation results showing the savings provided by buffer sharing, and the effect of buffer sharing on admission control.

All the analyses presented so far are based on a fixed reading order of streams within a cycle. [2, 3] explore the benefit of allowing the reading order to change from one period to another. The gain is a reduction in total seek time, whereas the price to pay may be a doubling of buffer requirements. In future work, we will study whether we can get the best of both worlds by integrating buffer sharing with variable reading orders.

4 Prefetching Strategies

4.1 Benefits of Prefetching

On receiving a new request for a stream (referred to as a new query from now on), the admission controller that we have discussed so far simply checks if there are enough disk bandwidth and buffers to satisfy the new query, using Equations 5 and 8. If there are enough resources, the query is activated. Otherwise, the query sits idle in the waiting queue. Consequently, there are resources – buffers and disk bandwidth – that are not utilized at all [2]. For instance, consider the situation mentioned in Example 1. If the disk bandwidth can support only 4 streams and there are 2MB buffering space, buffer sharing would render 1MB idle. In general, we measure the performance of our system by its throughput and the response time of queries. But given a system with pre-determined (at design time) disk bandwidth and amount of buffer space, the response time of queries are primarily determined by the utilization of disks and buffers. Thus, our goal here is to try to use these resources as much as possible. More specifically, in this section, we explore how data prefetching can maximize resources utilization, and thus lead to an increase in system throughput.

There are at least 3 ways that prefetching can help a query.

- First, if a query has a consumption rate P_i that is larger than R, then even after a query is activated (i.e. becoming one of the queries served in a cycle), the query cannot be consumed immediately without violating the continuity requirements. Thus, to reduce the time between activation and the beginning of consumption, a system can prefetch portion of this query while it is still waiting in the waiting queue.

- Second, even if a query S_{n+1} has a consumption rate P_i less than R, prefetching portion of this query before activation may reduce the response time of the query. To see that, let say that $S_1, \ldots, S_n$ are the activated queries. At some point, query S_1 has finished, and S_{n+1} is activated. For reasons apparent later in Section 5.1, the reading order may become $S_2, \ldots, S_{n+1}$. If no data has been prefetched for S_{n+1}, then S_{n+1} cannot be consumed until S_{n+1} starts reading, which is at the end of the cycle. However, if there is sufficient amount of prefetched data of S_{n+1}, consumption of S_{n+1} can start immediately at the beginning of the cycle. Thus, there is a difference in response time which may be as large as one cycle length.

- Third, prefetching portion of this query before activation has the effect of reducing the effective consumption rate of the query after activation. This is illustrated in Figure 3. The solid line represents the original consumption curve, whose slope is given by the consumption rate P_i. If an amount pf is prefetched, then the new, prefetched consumption rate is given by the slope of the dotted line. A simple analysis reveals that if T_i is the length of the query, the new, prefetched consumption rate is given by:

$$P_i^{pft} = P_i - \frac{pf}{T_i} \quad (11)$$

  Since the new rate is less than the original rate, there is a possibility that the new rate may pass the admission control test, while the old one may not. Whenever this happens, the response time of the query is substantially reduced (cf: Example 3 later).

In this section, we will first present a straightforward prefetching strategy SP. Then observing that the effectiveness of SP may be hindered by several shortcomings, we will develop another prefetching strategy IP which tries to maximize overall system throughput.

[2] In this paper, we only consider FIFO as the queueing discipline. It has the advantage of being simple and fair. Adopting other queueing discipline may require additional work to ensure fairness.

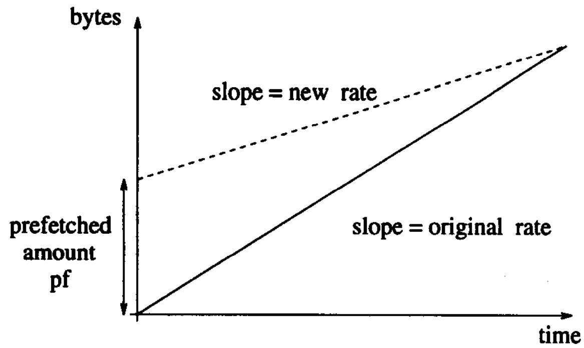

Figure 3: Reducing Consumption Rate by Prefetching

4.2 A Simple Prefetching Strategy: SP

Just like normal data retrieval from disk, prefetching requires both disk bandwidth and buffers. One obvious way to allow prefetching to happen is to dedicate a certain level of disk bandwidth and buffers to prefetching. But this would backfire as it reduces the disk bandwidth and buffers available to activated queries. Thus, we make sure that prefetching is not done at the expense of activated queries. To this end, recall that the cycle length t for the activated streams/queries $S_1, \ldots, S_n$ are bounded below and above respectively by Equations 5 and 8. If the system does not support prefetching at all, any value between the upper and lower bounds can be picked as the value of t. However, to support prefetching, an immediate question to answer is how to pick t so as to maximize prefetching, but not at the expense of the activated queries.

In fact, setting t to any value between the upper and lower bounds does not have any influence whatsoever on the completion times of the activated queries, as the completion time of a query is determined by its consumption rate and length [3]. Thus, as long as a value is picked between the lower and upper bounds, the activated queries will not be affected. Let us consider setting t to its lower bound. Then as discussed in Section 2, this corresponds to a disk utilization ρ of 1. In other words, all the disk bandwidth is used up for the activated queries, and nothing is left for prefetching. On the other hand, consider setting t to its upper bound. From the point of view of disk bandwidth allocation, this time there is ample room for prefetching because as discussed in Section 2, a longer cycle length corresponds to a lower disk utilization ρ. However, the trouble is that all the buffers are used up for the allocated queries. Thus, at the end, no prefetching can be done. Hence, the question to address is which value of t in between the upper and lower bounds maximizes prefetching.

There is actually another factor that affects the amount of prefetching that can be done. All the above analysis is based on the assumption that the cycle for the current collection of activated queries keep on going. Let T_{finish} denote the time the next activated query will have finished. The range bounding t is only valid before T_{finish}, after which the current cycle has to be changed anyway, and new calculations are required. Thus, the consideration of T_{finish} suggests a simple strategy (referred to as SP) to pick t so as to maximize prefetching. It equates the amount of data that can be retrieved in time T_{finish} with the amount of buffers that are available. This is formalized below. First, it is obvious that the amount of data that can be prefetched in time T_{finish} is: $D_{pf} = T_{finish} * R * (1-\rho)$. By substituting Equations 3 and 4, we get:

$$D_{pf} = T_{finish} * R * (1 - \frac{s}{t} - \frac{P}{R}) \quad (12)$$

On the other hand, according to Equation 7, the buffers available for prefetching is given by:

$$B_{pf} = B_{max} - \frac{t}{R} \sum_{i=1}^{n} P_i * (R - P_i) \quad (13)$$

To maximize prefetching, SP sets

$$D_{pf} = B_{pf} \quad (14)$$

This is a quadratic equation in t in the form of $at^2 + bt + c = 0$. Solving this quadratic equation in the standard way gives a positive solution v_0 (and a negative solution). If v_0 falls within the lower and upper bounds of t, which occurs more often than not, v_0 is the value of t. Otherwise, if v_0 is strictly less than the lower bound, t is set to the lower bound. And if v_0 is strictly greater than the upper bound, the upper bound becomes the value of t.

The equations presented above do not assume buffer sharing, and are based on Equation 7. Since prefetching is orthogonal to buffer sharing, a similar set of equations can be derived for the buffer sharing case based on Equation 10. Strategy SP is summarized in the following.

Strategy SP Let $S_1, \ldots, S_n$ be all the activated queries, as allowed by the admission controller. Let S_{n+1} be the query at the head of the waiting queue.

1. Use Equation 14 to determine the length t of the cycle for $S_1, \ldots, S_n$.
2. Use the remaining disk bandwidth and buffers to prefetch for S_{n+1} at the end of each cycle.
3. Prefetching stops when an activated query has finished, or the system has run out of buffers. □

[3] This is assuming normal termination, not preempted by such events as user quitting prematurely or system failures.

within cycle	t_1	...	t_n	remainder
read operations	S_1	...	S_n	idle

within cycle	t'_1	...	t'_n	remainder
read operations	S_1	...	S_n	prefetch S_{n+1}

Figure 4: Cyclic Activities with or without Prefetching

Figure 4 shows the disk activities within each cycle with or without prefetching. The first table shows the situation without prefetching. Within each cycle, the first t_1 seconds are spent reading S_1, the next t_2 for S_2, and so on. After all $S_1, \ldots, S_n$ have finished reading, the disk is idle for the remaining time of the cycle. The second table in Figure 4 shows the situation when prefetching strategy SP is used. It differs from the first table in two respects. First, its cycle length may be different from that without prefetching. Thus, the reading times for the activated queries are $t'_1, \ldots, t'_n$, instead of $t_1, \ldots, t_n$. More importantly, after the activated queries have finished reading, the disk may no longer be idle, and may be engaged in prefetching S_{n+1}.

4.3 Motivation for a More Intelligent Prefetching Strategy

Prefetching Strategy SP maximizes prefetching for the query S_{n+1} at the head of the waiting queue. Doing so, it may minimize the response time of S_{n+1}. However, as a result, S_{n+1} may use up too much system resources, particularly free buffers – for its own good, but not necessarily for the overall benefit of the system. More specifically, SP just lets S_{n+1} prefetch as much as possible, but does not consider whether S_{n+1} really needs that much data to get started once an activated query has finished. As shown in the example below, too much prefetched data only occupy buffer space, without doing any good to system performance.

Example 2 Consider a situation similar to the one described in Example 1. There are 4 activated queries, each with a consumption rate 240KB/s. And there is 1MB of buffer space left. Now consider a scenario where the query S_5 is the only query in the waiting queue with the same consumption rate. As discussed before, SP would allow S_5 to prefetch as much as possible, using up all 1MB of buffer space. However, as calculated in Example 1, 456KB is all that is needed for S_5 within a cycle. In other words, 456KB is sufficient to minimize the response time of S_5. Thus, the question is whether prefetching an extra 544KB can lead to any gain. The answer is no, because once the consumption of S_5 begins, its completion time depends entirely on its length and its consumption rate. Giving extra buffers does not help in any way. And in fact, it can be harmful to the entire system as there is now 544KB less of buffer space available. □

The above example suggests that while maximizing prefetching, the SP's approach of prefetching just for the query at the head of the waiting queue may not be sufficient. Thus, for a more effective prefetching strategy, the questions to be answered are: a) how to maximize prefetching, and b) how to determine how much to prefetch for a query in the waiting queue. The following example shows how looking ahead beyond the query S_{n+1} at the head of the waiting queue can help to determine the amount to prefetch for S_{n+1}.

Example 3 Consider the situation discussed in the previous example again. There are 4 activated queries with consumption rate 240KB/s each. Suppose there are now two queries in the waiting queue: S_5 and S_6 both with consumption rate 240KB/s. Further assume that the disk has a maximum reading rate of $R =$ 1150KB/s, and there is now 1.5MB of buffer space. Now let us consider the time when one of the activated queries has finished, and consider two different amounts of prefetched data of S_5.

First, assume that 456KB of S_5 has been prefetched, which would minimize the response time of S_5. By Equation 11, the new, prefetched consumption rate of S_5 is $240 - 456/30 = 225$, assuming that the total length of S_5 is 30 seconds. The question is whether S_5 and S_6 can be activated simultaneously. The answer is no because the total consumption rate $P = 3*240+225+240 = 1185 > 1150$.

Alternatively, assume that 1500KB of S_5 has been prefetched. Then, by Equation 11, the prefetched consumption rate of S_5 is $240 - 1500/30 = 190$. In this case, the total consumption rate $P = 3*240+190+240 = 1150$ which is $\leq R = 1150$. [4] Thus, as long as there are enough buffers to accommodate S_6, both S_5 and S_6 can be activated, reducing drastically the response time of S_6. Thus, the consumption rate of S_6 can be used to determine an appropriate amount to prefetch for S_5. □

The above example shows that prefetching S_5 for the appropriate amount can lead to a gain for S_5 and other queries in the waiting queue. It also leads to an interesting question: how to distribute prefetching among queries in the waiting queue. In other words, given the same amount of buffer space available for

[4] In practice, it is not so simple just to ensure that the total consumption rate is not greater than the maximum reading rate. As shown later in Strategy IP, what needs to be done is a full admission control test. But here we simplify the situation to illustrate the point that prefetching can lead to the activation of extra queries.

prefetching, how much of each query in the waiting queue should be prefetched so as to maximize the reduction in total consumption rate, thereby maximizing the number of queries that can be activated.

To answer this question, let us consider a "marginal gain" analysis on the buffers, quite similar to the one used in [7]. More specifically, for a query S_i, with an original consumption rate P_i, we calculate the reduction in consumption rate we would obtain if we prefetch one extra KB of S_i. By Equation 11, this value is equal to $P_i - P_i^{pft}$ which is equal to $\frac{1}{T_i}$. Thus, given queries S_1, S_2 whose lengths are T_1, T_2 respectively, prefetching more for the stream whose length is the shorter between T_1 and T_2 would result in a sharper drop in the combined consumption rate of the two streams. In other words, if the combined consumption rate has to drop below a certain value in order to pass admission control, prefetching more for the shorter query would require fewer buffers than prefetching for the longer one. Consider the following example.

Example 4 The previous example shows that in order to activate both S_5 and S_6 after one other query has finished, prefetching S_5 for 1500KB will do. Suppose the length of S_6 is 15 seconds. Then solving the equation $P = 3 * 240 + 240 + (240 - \frac{pf}{15}) = 1150$ indicates that if we prefetch S_6 entirely, only an amount pf of 750KB would be sufficient to activate both S_5 and S_6. Note that this amount is the bare minimum that allows both queries to be activated. If there are extra buffers, we can do more by prefetching one cycle of S_5 as well, so that not only are they activated, but both S_5 and S_6 can also be consumed immediately at the beginning of their first cycle. □

4.4 Prefetching Strategy IP

The prefetching strategy below, called IP which stands for "Intelligent Prefetching," finds the shortest query to prefetch, so as to maximize prefetching and the number of queries that can be activated once an active query has completed.

Strategy IP Let $S_1, \ldots, S_n$ be all the activated queries, as allowed by the admission controller. Among them, let S_j ($1 \leq j \leq n$) be the query that will finish the earliest. Also let $S_{n+1}, S_{n+2}, \ldots$ be the queries in the waiting queue, and B_{free} be the total number of buffers available to prefetching.

1. Use Equation 14 to determine the length t of the cycle for $S_1, \ldots, S_n$.
2. Initialize *target* to S_{n+1}, and *candidateSet* to S_{n+1} as well. Also set *finalAmt* to 0.
3. (** first chance **) If the combined consumption rate of all the streams in *candidateSet* is not greater than the consumption rate of S_j (i.e. $P_j \geq \sum_{S_k \in candidateSet} P_k$), go to Step 6.
4. (** second chance **) Otherwise,
   (a) Calculate the necessary prefetched consumption rate P_{target}^{pft} of *target* so that all the streams in *candidateSet* can possibly be activated when S_j has finished, i.e. $P_{target}^{pft} + \sum_{S_k \neq target; S_k \in candidateSet} P_k \leq P_j + (1 - \rho) * R$.
   (b) Use Equation 11 to calculate the amount that needs to be prefetched in order to reduce the consumption rate of *target* to P_{target}^{pft}, i.e. $targetAmt = (P_{target} - P_{target}^{pft}) * T_{target}$.
   (c) If $targetAmt > B_{free}$, then go to Step 5 to try the next condition.
   (d) Otherwise, use the admission control test given in Section 3.1 to determine if all streams in *candidateSet*, including the prefetched one, can get in a cycle with all the current activated queries except S_j. If the admission control test fails, go to Step 5.
   (e) Otherwise, set *finalTarget* to *target* and *finalAmt* to *targetAmt*. Go to Step 6.
5. (** third and final chance: both Steps 3 and 4 fail **)
   (a) Set *targetAmt* to B_{free}.
   (b) Use Equation 11 to calculate the prefetched consumption rate P_{target}^{pft} of *target*, i.e. $P_{target}^{pft} = P_{target} - \frac{targetAmt}{T_{target}}$.
   (c) Use the admission control test given in Section 3.1 to determine if all streams in *candidateSet*, including the prefetched one, can get in a cycle with all the current activated queries except S_j. If the admission control test fails, go to Step 7.
   (d) Otherwise, set *finalTarget* to *target* and *finalAmt* to *targetAmt*. Go to Step 6.
6. (** try to see if more queries can be activated **) Consider the next query S_{next} in the waiting queue that is not in *candidateSet*. Add S_{next} to *candidateSet*. Compare the length of S_{next} with the length of *target*. Set *target* to be the stream with the shorter length. Go back to Step 3.
7. (** no more queries can be activated **) If $finalAmt > 0$, prefetch *finalTarget* for the amount *finalAmt*. □

In the above strategy, the purpose of *candidateSet* is to ensure FIFO in the activation of queries, even though as argued in the "marginal gain" analysis above, it is possible to prefetch S_{k+1} without prefetching S_k. In each iteration of IP, the stream with the shortest length in *candidateSet* is chosen to be the *target* stream for possible eventual prefetching. Then there are three possibilities for all the queries in the *candidateSet* to be activated, once S_j has completed (i.e. the next activated query to finish). The first case is when the combined consumption rate of all those in *candidateSet* does not exceed the consumption rate of S_j. In this case, all queries in *candidateSet* are guaranteed to be activated once S_j has completed. In addition, nothing needs to be prefetched in this case [5]. Execution then goes to Step 6 to try to see if more queries in the waiting queue can be activated. A new *target* is found, and a new iteration begins.

If the first condition fails in Step 3, execution goes to Step 4 to see if the second possibility would work out. In this case, IP tests if a sufficient amount of *target* can be prefetched so that all queries in *candidateSet* can be activated, provided that this amount does not exceed the number of buffers currently available to prefetching (cf. Step 4c). If admission control in Step 4d verifies that all queries can be activated with the help of prefetching, both *target* and the prefetching amount *targetAmt* are recorded in the variables *finalTarget* and *finalAmt*. Execution then goes to Step 6 to try to add another query from the waiting queue to *candidateSet*, and a new iteration begins.

If both the conditions in Steps 3 and 4 fail, IP tries the "last resort." It simply tests to see if using all free buffers to prefetch for *target* will be sufficient to activate all queries in *candidateSet*. If admission control returns a positive answer, all the necessary operations will be taken in Step 5d and 6, and a new iteration begins.

If all three conditions in Steps 3, 4 and 5 fail, it is an indication that not all queries in *candidateSet* can be activated. More precisely, all but the last added query in *candidateSet* can be activated once S_j has completed. Step 7 prepares for this event by prefetching *finalTarget* for the amount *finalAmt*. As shown in Figure 4, prefetching occurs at the end of each cycle.

Notice that as presented above, IP is only concerned with maximizing the number of queries that can be activated. As discussed in the previous example, IP can easily include a Step 8 that would prefetch one cycle worth of data for each query that would be activated, so that every one can be consumed immediately at the beginning of the first cycle. Furthermore, in the case when no query in the waiting queue can be activated even after S_j has completed (i.e. S_{n+1} is the only query in *candidateSet*), another thing Step 8 could do is to use SP to prefetch as much as possible for S_{n+1}. This would take care of the situation when the consumption rate of S_{n+1} needs to be substantially reduced before S_{n+1} can be activated. Last but not least, the admission control used in IP above does not consider buffer sharing. Equation 10 can be used in the place of Equation 7 (and thus Equation 8) in admission control, if buffer sharing is used.

[5] This is the case as far as query activation is concerned. But if there are enough buffers available at the end, queries in *candidateSet* may be prefetched so that they can be consumed immediately at the beginning of their first cycle.

Example 5 Let us apply Strategy IP to the situation discussed in the previous example. Let us assume that S_1 is the activated query that will finish the earliest. In the first iteration of IP, S_5 alone is considered in Step 3. Since S_5 has the same consumption rate as S_1, S_5 can certainly take the place of S_1 and be activated once S_1 has completed. Thus, execution goes to Step 6, in which S_6 is added to *candidateSet*. Since S_6's length is shorter than S_5's, S_6 becomes the new *target*.

In the next iteration of IP, obviously Step 3 fails. Now based on the calculations given in the previous example, the prefetched consumption rate of S_6 is P^{pft}_{target} = 190KB/s, and the prefetched amount is *targetAmt* = 750KB. Assuming that the admission control test in Step 4d is passed, *finalTarget* is set to S_6 and *finalAmt* to 750KB. Then in Step 6, another query S_7 is added from the waiting queue to *candidateSet*, and a new iteration begins.

Suppose S_7 has the same rate and length as S_6, and is the new *target*. It is not difficult to verify that Steps 3, 4 and 5 fail in this iteration. Thus, execution goes to Step 7, and the final decision is that S_6, which is *finalTarget*, will be prefetched for 750KB. As discussed before, if there is a Step 8 in IP to minimize response time, S_5 will also be prefetched so that the consumption of S_5 can start immediately at the beginning of its first cycle. □

This concludes the presentation of our prefetching strategies. Next we will show preliminary simulation results evaluating the effectiveness of IP and SP, as well as buffer sharing.

5 Preliminary Simulation Results

5.1 Details of Simulation Package

We have implemented a discrete-event simulation package to evaluate the techniques proposed in this paper. The package runs under Unix on Sparc-stations, and consists of about 5,000 lines of C code. For ease of coding, all the queries to be executed in a simulation are submitted to the waiting queue at the beginning of the simulation. Thus, the main outputs of the simulation package do not include response times of queries, but include such statistics as peak and average disk and buffer utilizations, and the total time to complete all queries. Furthermore, to make our simulations as close to reality as possible, we have implemented the following features in our package.

- As observed in [9], a transient period is required before a new Stream S_{n+1} can be added to a (new) cycle. This is because the new cycle length t' (for one more stream) is strictly larger than the current cycle length t. Thus, if we directly serve S_{n+1} at the end of the current cycle, starvation will occur for all the queries $S_1, \ldots, S_n$ in the current cycle, because they only have data buffered for a cycle of length $t < t'$. Apart from the steady states, our package also simulates the transient period. For more information, see [8].
- When an activated query has completed, there are two ways to invoke the admission controller. One way is to wait till that particular cycle ends; the other is to wake up the controller immediately after the query has finished (even amidst a cycle). The former policy, while much easier to implement, does not optimize system performance, especially when the cycle length is long and the disk utility is low. We have implemented the latter policy, and found out that system performance is improved.

Apart from making our simulation package as close to reality as possible, we have designed and run our simulations based on real figures (e.g. minimum and maximum seek times equal to 5 and 25ms respectively). We will give further details on all the simulations presented below.

5.2 Effectiveness of Buffer Sharing

In Section 3.2, we have analyzed that buffer sharing can lead to a 50% reduction in total buffer requirement, when the disk utilization ρ is equal to 1. Here we simulated a situation when ρ keeps changing and has an average value less than 1. In this series of simulation, we used 50 queries, each with consumption rate 240KB/s. The lengths of the queries were from 20 to 120 seconds, with the average being 60 seconds. In order to support a sufficiently high number of concurrent queries, the maximum disk reading rate was set to R = 2000KB/s. The graph in Figure 5 shows the minimum buffer space needed, when the number of concurrent queries varies from 3 to 7 – with and without buffer sharing. As expected, in all cases, buffer sharing requires less buffer space than without buffer sharing. The savings in buffer space was between 20% to 40%, depending on the average disk utilization.

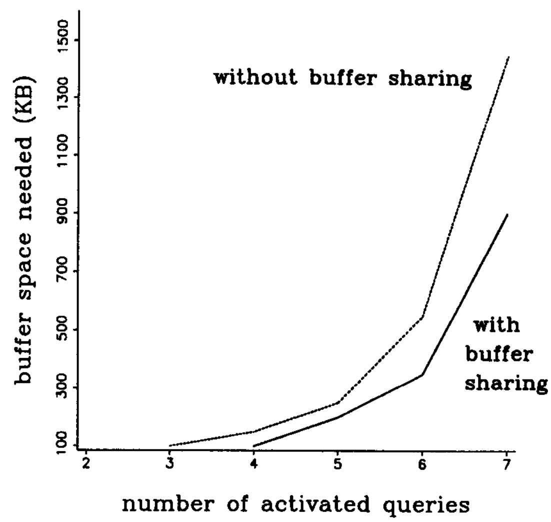

Figure 5: Benefit of Buffer Sharing

5.3 Effectiveness of Prefetching Strategy IP

In this series of simulation, we evaluated the effectiveness of our prefetching strategies. We again used 50 queries, each with consumption rate 240KB/s, and length 90 seconds. The maximum disk reading rate was set to 1000KB/s. The graphs in Figures 6 and 7 show the time taken to complete the 50 queries and the average disk utilization with varying amounts of buffer space. In both graphs, the x-axis is the amount of buffer space, varying from 5MB to 8.5MB. In Figure 6, the y-axis is the total time taken to complete 50 queries using IP and SP, normalized by the time taken without prefetching. Thus, the horizontal line at 1.0 in Figure 6 represents the situation without prefetching. With small amounts of space available to prefetching, IP does not lead to any gain in performance. However, as more and more space becomes available, IP is able to activate more and more queries faster than if no prefetching is allowed. Consequently, the total time taken becomes smaller. As shown in Figure 6, IP could lead to a 30% savings in total time taken. Alternatively, the throughput of a system using IP could be 3/7 ≈ 40% higher.

The performance gain caused by IP can be best explained by the graph in Figure 7. If no prefetching takes place, the average disk utilization is around 0.8. But as more buffer space becomes available to prefetching, IP is able to better utilize the disk by prefetching, and the average disk utilization gradu-

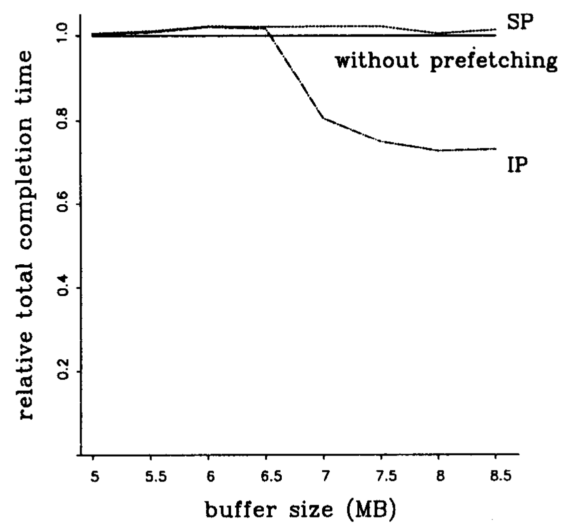

Figure 6: IP vs SP: relative total time taken

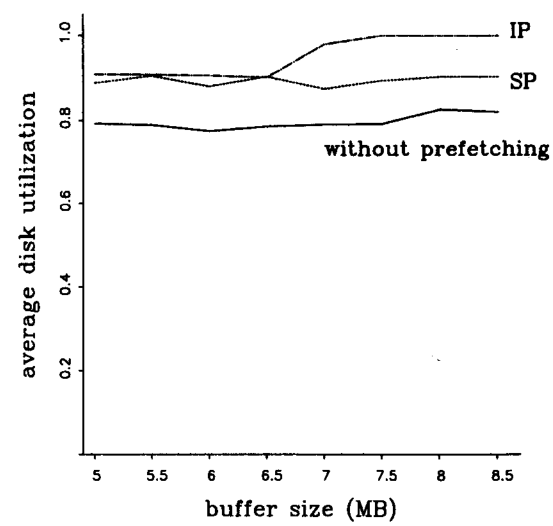

Figure 7: IP vs SP: average disk utilization

ally climbs up to 1.0. Moreover, the utilization of buffers follows a similar trend. Another interesting thing shown in Figure 7 is that the average disk utilization for SP is still higher than if no prefetching is allowed. This is an indication that while the disk is kept busy by prefetching, the way that SP conducts prefetching is problematic and totally ineffective.

The series of simulation discussed above did not allow buffers to be shared. In another series of simulation, we allowed buffers to be shared, and used the version of admission control that is based on Equation 10, but not on Equation 7. The results of this series of simulation were very similar to those presented above. The only difference was that buffer sharing saved a few hundred KBs of buffer space, and made it available to prefetching. Thus, the point when IP started to show improvement now began a few hundred KBs earlier than was shown in Figure 6.

5.4 Discussions: Applicability of Prefetching to General Multimedia Systems

Our preliminary simulation results indicate that appropriate prefetching can lead to increase in throughput, disk utilization and buffer utilization. However, in order to have higher throughputs and thus lower response times of queries, the price to pay is certainly availability of buffer space. As shown in our examples and simulation results, we believe that the price is not high – provided that the streams are short, say below 5 minutes in length. As far as news on-demand systems are concerned, a large class of news clips falls within this range. However, a natural question to ask is whether prefetching has a role to play in other multimedia systems.

Consider multimedia database management systems. We believe that prefetching indeed has a major role to play in tuning the performance of such systems. This is because for a large class of applications, the audio and video components tend to be short. For example, for applications such as the one described in [5], audio and video may not be the only media, and may work hand-in-hand with other media such as text and images. Audio and video components may also play the role of annotations or illustrations. Moreover, many applications may require frequent user interaction.

What about the other extreme: movies on-demand systems? Unlike those cases discussed above, movies on-demand is concerned with supplying video and audio data to users for long durations and with relatively little user interaction. By Equation 11, reducing the consumption rate of a movie by just 1KB/s requires T KB buffer space, where T is the length of the query in seconds. For example, if a movie is 90 minutes long, this amount of buffer space is already 5.4MB. And to reduce the consumption rate by 50KB/s (as in Example 5), 270MB of buffer space is needed! As shown in Equation 11, the amount of buffer space needed for prefetching (and IP) to work is linearly proportional to the length of the movie. However, on the positive side, consider the benefit of prefetching. Recall that prefetching has the effect of activating as many queries (movies) as possible. If prefetching is not used, and a movie M_0 cannot be activated immediately, it has to wait for an activated movie M_1 to finish. Thus, the waiting time of M_0 depends linearly on the length of M_1. In other words, if prefetching is the difference between whether a movie can or cannot be activated immediately, the difference in response time, like the amount of buffer space needed, is linearly proportional to the length of the movies. As an example, this difference in response time may be 30 minutes. Thus, while long queries magnify the buffer space needed for prefetching to work, they also magnify the benefits of prefetching. It is certainly up to an enterprise to decide which is more important and costly: 270MB or 30 minutes.

6 Conclusions

Providing effective multimedia support in database management systems is a topic of great interest and value. In this paper, we consider one of the key problems encountered in such systems. Given a fixed amount of buffer space and disk bandwidth both predetermined at design time, we study how to maximize the throughput of the system. Our approach is to maximize the utilizations of buffers and disk. To achieve this goal, we have first proposed a buffer sharing scheme. Analysis and simulation results indicate that buffer sharing could reduce total buffer consumption by as much as 50%. Second, we have developed the prefetching strategy IP which aims to maximize prefetching and the number of queries that can be activated. Preliminary simulation results show that IP could be quite effective in maximizing the effective use of buffers and disk, and could lead to a 40% increase in system throughput. Finally, as argued in Section 5.4, we believe that the proposed techniques can also be valuable to multimedia systems other than news on-demand systems and database management systems.

In ongoing work, we are studying how to implement the proposed techniques in a distributed continuous-media file system. Key issues to be addressed include how to extend the proposed techniques to support multiple disks and network buffering, and how to effectively implement prefetching and buffer sharing, when the reading orders from one cycle to the next can or cannot be changed.

Acknowledgements

Research partially sponsored by NSERC Grants OGP0138055 and STR0134419, IRIS-2 Grants HMI-5 and IC-5, and CITR Grant on "Distributed Continuous-Media File System."

References

[1] D. Anderson, Y. Osawa and R. Govindan. (1992) *A File System for Continuous Media*, ACM Trans. on Computer Systems, 10, 4.

[2] M. Chen, D. Kandlur and P. Yu. (1993) *Optimization of the Grouped Sweeping Scheduling with Heterogeneous Multimedia Streams*, Proc. ACM-Multimedia, pp 235–242.

[3] J. Gemmell. (1993) *Multimedia Network File Servers: Multi-channel Delay Sensitive Data Retrieval*, Proc. ACM-Multimedia, pp 243–250.

[4] J. Gemmell and S. Christodoulakis. (1992) *Principles of Delay-Sensitive Multimedia Data Storage and Retrieval*, ACM Trans. on Information Systems, 10, 1, pp 51–90.

[5] R. Goldman-Segall. (1990) *Learning Constellations: a Multimedia Research Environment for Exploring Children's Theory-Making*, Constructionist Learning, ed. I. Harel, Cambridge, MA, MIT Media Laboratory.

[6] G. Miller, G. Baber and M. Gilliland. (1993) *News On-Demand for Multimedia Networks*, Proc. ACM-Multimedia, pp 383–392.

[7] R. Ng, C. Faloutsos and T. Sellis. (1991) *Flexible Buffer Allocation Based on Marginal Gains*, Proc. ACM-SIGMOD, pp 387–396.

[8] R. Ng and J. Yang. (1994) *Maximizing Buffer and Disk Utilizations for Multimedia Systems*, Technical Report, University of British Columbia.

[9] P. Venkat Rangan and H. Vin. (1991) *Designing File Systems for Digital Video and Audio*, Proc. ACM Symposium on Operating Systems Principles, pp 69–79.

[10] A. Reddy and J. Wyllie. (1993) *Disk Scheduling in a Multimedia I/O System*, Proc. ACM-Multimedia, pp 225–233.

[11] L. Rowe and B. Smith. (1992) *A Continuous Media Player*, Proc. 3rd Intl. Workshop on Network and OS Support for Digital Audio and Video.

[12] K. Tindell and A. Burns. (1993) *Scheduling Hard Real-Time Multimedia Disk Traffic*, Technical Report, University of York, England.

[13] C. Yu, W. Sun, D. Bitton, Q. Yang and R. Bruno. (1989) *Efficient Placement of Audio Data on Optical Disks for Real-Time Applications*, Communications of ACM, 32, 7, pp 862–871.