

Zone-Bit-Recording-Enhanced Video Data Layout Strategies

Shenze Chen, Manu Thapar
Broadband Information Systems Laboratory
HPL-95-124
November, 1995

video-on-demand,
VOD, video layout,
zone-bit-recording,
disk I/O performance

Rapid progress in high speed networking and mass storage technologies has made it possible to provide *video-on-demand* (VOD) services, that deliver movies to viewers' homes on request. In this paper, we study video data layout issues in the video server design. Specifically, we present a family of novel video data layout strategies, called *Zone-Bit-Recording-Enhanced* (ZBRE) layout schemes, which take into account the multiple-zone-recording feature of modern disk drives. The ZBRE layout schemes can be applied to either individual disks or disk arrays, such as RAID 3. Simulation results show that, by carefully laying out popular movies in the outer zones and aggregating "hot" movies together, disk performance can be improved significantly, with an up to 23% higher throughput than by randomly laying out data on disks. More importantly, these performance benefits are obtained without incurring any extra hardware cost.

Internal Accession Date Only

To be published in and presented at *International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS '96)*, San Jose, California, February 1-3, 1996.

© Copyright Hewlett-Packard Company 1995

1. Introduction

Rapid progress in high speed networking and mass storage technologies has made it possible to provide *video-on-demand* (VOD) services, that deliver movies to viewers' homes on request. This technology requires three major components: *video servers*, *distribution networks*, and *set-top boxes*. Typically, digitized videos are compressed and stored in video servers to which customers send viewing requests. Upon receiving these requests, the video servers retrieve the movies from mass storage and deliver them over networks to the customers' homes [1]. At the customers' homes the movie streams are decompressed by set-top boxes and played back on TV screens. In order to guarantee the quality of these video-on-demand services, movie streams must continuously flow from the mass storage in servers to the set-top boxes in customers' homes.

In this paper, we focus on video data layout issues in the video server design. The movie layout schemes proposed here can be applied to either individual disks or disk arrays, such as RAID 3. Typically, movies are stored in multiple disks or RAIDs [2] that are attached to the video servers. Since in multimedia and video applications, a large amount of data needs to be retrieved from disks in a timely fashion, and disks account for most system costs, the disk I/O subsystem needs to be carefully designed to meet both performance and cost goals. The multiple-zone-recording technique has been used by disk manufacturers for years to increase disk capacities. This feature, however, has not been widely utilized by system designers to improve performance. For example, in most UNIX[†] systems, data is randomly placed on disks. In this study, we explore the potential benefit achievable by utilizing the disk multiple-zone-recording feature for video and multimedia applications. Specifically, we present several novel video data layout strategies, called *Zone-Bit-Recording-Enhanced (ZBRE)* layout schemes. The idea behind the ZBRE layout schemes is very simple: placing the most popular movies in disk outer zones and aggregating these "hot" movies together to take advantage of high bandwidth in outer zones. Our simulation results show that, by using these layout strategies, disk performance can be improved significantly, with an up to 23% higher throughput than by randomly laying out the data. We also study the trade-off between laying out hot movies in the outer zones and laying out hot movies in the center zones. Placing hot movies in outer zones can potentially improve bandwidth, and placing hot movies in center zones can potentially improve seek time. The results show that the gain from reducing seek time does not offset the loss of the high data bandwidth in the outer zones. Under our proposed ZBRE schemes, the disk throughput can be very close to that of the fastest tracks. More importantly, these performance benefits are obtained without incurring any extra hardware cost. Since storage costs are a major component of video servers, these improvements have direct impact on reducing system costs.

Video server I/O subsystem design is a relatively new area for which little can be found in the literature. Rangan *et al.* [3,4,5] presented a data layout scheme that divides a movie into blocks. These blocks are scattered on to disk surfaces, with "gaps" between adjacent blocks. Keeton and

[†] UNIX is a trade mark of AT&T Bell Labs.

Katz [6] presented an interesting scheme to layout video data that can be played back in different resolutions. This scheme requires video compression algorithms that support this decomposition by resolution. Berson *et al* [7] proposed a staggered method of striping multimedia information across multiple disks based on worst case analysis. Kenchammana-Hosekote and Srivastava [8] studied the issue of scheduling storage devices to guarantee rate requirements for continuous media. This work is also based on worst case analysis. Recently, Birk [9] proposed a scheme that pairs each track in an outer zone with a track in an inner zone and places the movies in these track pairs. His results show that, with this track-pairing, the throughput is approximately 20% lower than that of the fastest tracks and 40% higher than that of the innermost tracks. Other work on video server buffer management and multicasting can be found in [10] and [11].

The remainder of the paper is organized as follows: Section 2 describes the multi-zone structure of modern disks; Section 3 presents the ZBRE movie layout schemes; a performance evaluation of different layout strategies is given in Section 4; and Section 5 summarizes this paper.

2. Multi-Zone Structures of Modern Disk Drives

Typically, magnetic disks consist of several platters, with each platter containing many tracks. Older disks used a fixed angular density scheme, which records a fixed number of bits per track. Since the perimeters of inner tracks are smaller than the perimeters of outer tracks, the bit density of outer tracks is a lot less than that of inner tracks. In order to achieve higher capacity, modern disks adopt a fixed linear density scheme that records a fixed number of bits per unit length within each track. To simplify the management of disk transfers, disk manufacturers actually use a multi-zone structure to approximate the fixed linear density scheme that divides tracks on each platter into several zones. All tracks within the same zone record the same number of bits per track, but tracks in outer zones record more bits than those inner zones.

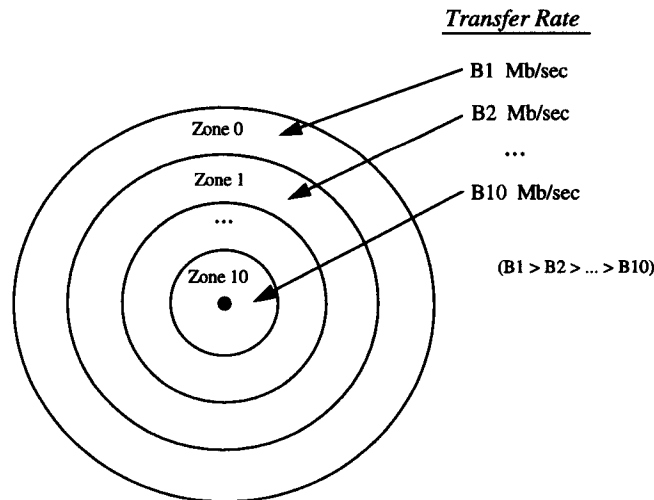


Figure 1: Multi-Zone Structure of Modern Disks (HP C2490A).

Figure 1 shows the multi-zone structure of the HP C2490A 3.5" drive, which consists of 11 zones [12]. The detailed structure is listed in Table 1. Other disks studied in this paper include the HP C2247 3.5" drive which consists of 8 zones, and the Seagate ST32550 3.5" drive which includes more zones per surface.

From Table 1, we see that the ratio of data sectors for the outermost and innermost zones is 1.6. This implies that, under the same rotational speed, the outermost zone data bandwidth will be 1.6 times higher than that of the innermost zone when transferring data between the disk media and the controller. The ratios for the HP C2247 and the Seagate ST32550 drives are 1.7 and 1.5, respectively. For some other drives, such as the DEC DSP3000L 3.5" drive, this ratio can be as high as 2.0.

Table 1: Zone Structure of HP C2490A Drive.

	Tracks per Surface	Sectors per Track (512B/sector)
zone 0 (outer)	174	108
zone 1	310	104
zone 2	280	100
zone 3	308	96
zone 4	256	92
zone 5	250	88
zone 6	236	84
zone 7	218	80
zone 8	204	76
zone 9	218	72
zone 10 (inner)	128	68

3. Video Data Layout Strategies

In this section, we describe a family of ZBRE video data layout strategies, which take movie "hotness" or "temperature" into account when placing video data on disks.

A large video server stores a mix of "hot" and "cold" movies, with the more recent, popular movies being hot, and older, less frequently watched movies being cold. The temperature of a movie is a measure of its bandwidth requirement and is a function of the number of estimated simultaneous viewers (based on release date, box office income and review rating), stream data rate, and usage history.

$$\text{temperature} = f(\text{release_date}, \text{box_income}, \text{review_rating}, \text{data_rate}, \text{usage_history})$$

Typically, a two-hour movie needs 2 to 3 GB of storage space, depending on the coding bit-rate. Therefore, multiple disks are needed in order to store multiple movies. In our video server model,

there are n movies stored on k disks attached to the same server, with each movie striped across all of the k disks (or arrays) to avoid a single disk becoming a bottleneck, as shown in Figure 2. The n movies are ranked according to their temperatures, M_1, M_2, \dots, M_n , with M_1 the *hottest* and M_n the *coldest* movies. When striping is done, each movie is placed on the same cylinder for each disk, depending on the layout policy used. Since each disk keeps a portion of all the n hot and cold movies, each disk faces the same I/O workload and access pattern. The server is expected to support multiple concurrent video streams, and therefore a disk is expected to move its arm when switching from stream to stream.




Figure 2: All Movies Striped across all Disks.

We define a *percentile-distance*, d , of a zone to be its percentile position relative to the outermost zone. For example, in Figure 1, zone 5 (in the center) has a percentile-distance of 0.5; zone 10 (the innermost zone) has a percentile-distance of 1; and zone 0 (the outermost zone) has a percentile-distance of 0.

We consider the following video data layout strategies:

- **ZBRE.0 Strategy**

places hot movies in outer zones and cold movies in inner zones, starting with the hottest movie M_1 in the outermost zone 0. This strategy takes advantage of the high bandwidth of the outer zones and directs most of the I/O requests to these outer zones. Within each zone, a movie can be allocated to any contiguous cylinders.

- **ZBRE.1 Strategy**

is considered an enhancement to ZBRE.0. While ZBRE.0 may randomly place movies within each zone, ZBRE.1 places the hottest movie M_1 on the innermost tracks of zone 0, followed by M_2 , and so on until zone 0 is full. Then the next hottest movie is placed in zone 1 starting from its outermost track. This pattern is repeated for each zone until the whole disk is filled. This strategy keeps the high bandwidth of the outer zones and reduces the seek time by placing the hottest movie toward the central portion of the disk and by aggregating together other hot movies as much as possible.

- **ZBRE.s Strategy**

places the hot movies starting at a sub-outer-zone which has a percentile-distance of $d > 0$, where d is an adjustable parameter. In the case of the HP C2490A drive, as shown in Figure 1, if the percentile-distance $d = 0.2$ and each zone holds two movies, then the ZBRE.s strategy lays out the hottest movie M_1 at the innermost track of zone 2. When zone 2 is filled, the next hottest M_3 goes to zone 1. Zones are filled in the order: zone 2, zone 1, zone 3, zone 0, zone 4, zone 5, zone 6, ... , zone 10, as shown Table 2.

Table 2: Layout of Movies for ZBRE.s.

zone 0	zone 1	zone 2	zone 3	zone 4	zone 5	zone 6	zone 7	zone 8	zone 9	zone 10
$M_8 M_7$	$M_4 M_3$	$M_2 M_1$	$M_5 M_6$	$M_9 M_{10}$	$M_{11} M_{12}$	$M_{13} M_{14}$	$M_{15} M_{16}$	$M_{17} M_{18}$	$M_{19} M_{20}$	$M_{21} M_{22}$

By placing hot movies toward the center of the disk, seek time can potentially be reduced, but some bandwidth for the outermost zone is sacrificed when ZBRE.s is compared to ZBRE.0 and ZBRE.1. Therefore, the ZBRE.s strategy performs better if the gain due to reduced seek time is greater than the loss due to reduced bandwidth. We will further examine this trade-off in Section 4.2. It is clear that the percentile-distance d used in this strategy should not be greater than 0.5 (i.e., the center of the disk), since placing hot movies in the inner zones doesn't gain anything.

- **ZBRE.r Strategy**

places the hottest movies in the outermost zone 0 in the same way as ZBRE.0, i.e., those movies can be placed in any contiguous cylinders within zone 0. This strategy differs from ZBRE.0 in that it places the rest of the movies randomly in the remaining zones. The performance of ZBRE.r is worse than that for other members of the ZBRE family. This strategy, however, makes it easy to replace hot movies on disks. For example, to store a new hot movie, we only need to make room for it in zone 0. We do not need to follow the layout patterns imposed by the other ZBRE strategies. Therefore, this strategy trades some performance for ease of implementation. We will study this performance trade-off later in Section 4.4.

- **RANDOM Strategy**

randomly allocates movies on disks, but each movie still occupies contiguous cylinders. The RANDOM strategy is provided as a basis for our performance studies.

4. Performance Evaluation

In this section, we study the performance of the video data layout strategies by simulating the three fast/wide SCSI disk drives mentioned in Section 2. In order to quantitatively characterize the temperature of movies, we use the "sth degree of skew" function, which defines the probability p_m of a customer request going to movie M_m .

$$p_m = \frac{m^s}{\sum_{i=1}^{i=n} i^s} \quad m = 1, 2, \dots, n. \quad (1)$$

where n is the total number of movies stored in the video server. Notice that, when parameter $s = 0$, the movie temperature is uniformly distributed, i.e., each movie has an equal chance of being selected for viewing. When $s = -1$, it is the well known Zipf's distribution [13], in which about 49% of viewers select the top 10% of hot movies. As the value of parameter s decreases, the viewers' selections become more skewed, which is likely to happen in practice based on current information about usage patterns. For example, when $s = -2$, about 88% of the viewers select the top 10% of the hot movies. In Table 3, we list the skew degree s and its corresponding $\alpha - \beta$ distribution, by which α percent of viewers select the top β percent of the hot movies.

Table 3: Movie Skew Distribution.

Skew Degree	$\alpha - \beta$ Distribution (α % of requests to β % of movies)
0	10% - 10%
- 0.5	25% - 10%
- 1.0	49% - 10%
- 1.5	73% - 10%
- 2.0	88% - 10%
- 2.5	95% - 10%
- 3.0	98% - 10%

In our simulation study, 40 movies are laid out on disks according to the strategies described in the previous section. Each movie requires an average of 2GB capacity. For each run, more than 100,000 I/O requests are generated, and each request is directed to movie M_m with probability p_m as defined by the hotness distribution in Equation (1). If a movie occupies l cylinders, then a request randomly goes to one of the l cylinders. The disk parameters are listed in Table 4. While the capacities of these drives are relatively small (1-2GB), disk vendors can easily achieve a higher capacity by adding more platters and maintain the same geometry and performance parameters.

Table 4: Disk Parameters.

	Seagate ST32550	HP C2490A	HP C2247
Avg. Seek (ms)	8.5	9.1	10.7
RPM	7200	6400	5400
Xfer Rate (Mb/sec)	47.5 - 72	36 - 57	24.8 - 42.4

Confidence intervals of 95% are obtained by using the method of independent replications. The confidence interval widths are less than 2% of the point estimates of the disk service times (these intervals are not shown in our figures in order to clearly show the results).

4.1. Performance of Various Layout Strategies

In this section, we compare the performance of the video data layout strategies presented in Section 3. We varied the I/O transfer size from 64KB to 256KB and observed the same performance ordering amongst these layout strategies. For ease of discussion, we only show the results for an I/O transfer size of 128KB. We initially set the movie hotness skew degree to -2 and examine the effects of varying skew degrees in Section 4.3.

First, we consider the three variations of the ZBRE family: ZBRE.0, ZBRE.1, and ZBRE.s (the ZBRE.r strategy will be discussed in Section 4.4). For the sub-outer-zone strategy ZBRE.s, we used the optimal number for the parameter d as discussed in Section 4.2. In general, their performance is similar, with one strategy slightly better than others. In particular, ZBRE.1 is consistently better than ZBRE.0, since ZBRE.1 performs seek optimizations within zones. Between the sub-outer-zone strategy ZBRE.s and outer-zone strategy ZBRE.1, the performance orderings vary according to the movie hotness skews, as will be discussed in Section 4.2.

Table 5: Performance of ZBRE Strategies.

	ZBRE.0	ZBRE.1	ZBRE.s
Seek Time (ms)	3.7054	3.3849	3.1906
Rotational Latency (ms)	4.1702	4.1702	4.1702
Transfer Time (ms)	17.2870	17.2870	17.5354
Total	25.1625	24.8421	24.8965

Table 5 shows the performance of three ZBRE strategies for the Seagate ST32550 drive, with a movie hotness skew degree of -2. The time is the mean time for 128KB I/O transfers. Clearly, ZBRE.1 has the same rotational latency and transfer time as ZBRE.0, but less seek time. On the other hand, ZBRE.s has less seek time than ZBRE.1, but a higher transfer time. In the case of skew degree $s = -2$, in which about 88% of customers select the top 10% of the hot movies, the overall performance of ZBRE.1 is slightly better than ZBRE.s. This is also true for both the HP C2490 and C2247 drives.

In Figure 3 we plot the performance improvement achieved by using the proposed ZBRE video data layout strategies as compared to randomly laying out movies on disks. Figure 3 (a) shows the mean disk service times under the RANDOM and ZBRE.1 strategies for the three disks used in this study. Figure 3 (b) gives the corresponding disk throughput. The results show that ZBRE.1 performs 20% better than RANDOM for the Seagate ST32550 and HP C2490 drives. For the HP C2247 drive, ZBRE.1 outperforms the RANDOM strategy by 23%.

When the transfer size is decreased to 64K, the improvement for the HP C2490A and Seagate ST32550 drives increases to 23%, and when the transfer size is increased to 256K, the improvement decreases to around 17%. For the HP C2247 drive, ZBRE.1 constantly outperforms the RANDOM strategy by 23% for all tested transfer sizes.

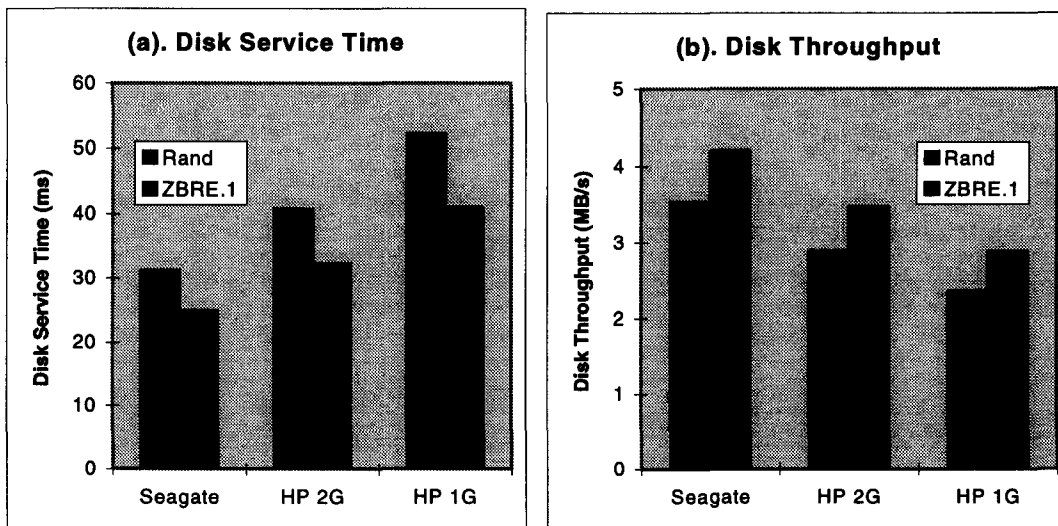


Figure 3: Comparison of ZBRE and RANDOM Layout Strategies.

Finally, in the case of skew degree $s = -2$, the data bandwidth achieved by ZBRE.1 is approximately 98% of the outermost zone bandwidth for the Seagate ST32550 and HP C2490A drives, and is 99% for the HP C2247 drive. When the movie skew decreases to -1, the achievable data bandwidths for ZBRE.1 are 95%, 91%, and 94% for the Seagate ST32550, HP C2490A, and HP C2247 drives, respectively.

4.2. Varying the Parameter d under ZBRE.s Strategy

In this subsection, we examine the effect of varying the percentile-distance parameter d under strategy ZBRE.s, which places the hottest movie on sub-outer zones, instead of the outermost one. When $d=0.5$, the hot movies are placed in the center of the disks.

Table 6: Disk Service Time (ms) under ZBRE.s (HP C2490A).

Percentil	Skew Degree = -1			Skew Degree = -2		
	64K	128K	256K	64K	128K	256K
0	25.24	37.28	61.37	20.91	32.22	54.85
0.1	25.04	37.22	61.59	21.29	32.95	56.27
0.2	25.11	37.54	62.39	21.79	33.88	58.05
0.3	25.41	38.16	63.65	22.36	34.93	60.07
0.4	25.78	38.90	65.13	22.89	35.99	62.17
0.5	26.32	39.86	66.95	23.47	37.15	64.50

Table 6 shows the mean disk service time, which includes seek time, rotational latency, and the data transfer time from the media to controller, under different values of percentile-distance d for

the HP C2490A drive. When the movie hotness skew degree is -1 (Zipf's law), we observe that the best disk performance is achieved when $d=0.1$, except when the transfer size is 256K. When the hottest movies are placed further towards the center of the platters, the mean disk service time increases, which means that the gain from reducing seek time does not offset the loss of the high data bandwidth in the outer zones.

When movie temperature becomes more skewed (skew degree $s = -2$), placing the hottest movies in the outermost zone results in the best performance, as shown in Table 6. When movie hotness is highly skewed, more I/O requests go to those hot movies, thus requiring fewer seeks. Therefore, the savings on seek times achieved by placing the hot movies in the sub-outer zones decreases.

The same behavior has been observed with Seagate ST32550 drives, with the exception that, when the transfer size is equal to 256K, the Seagate drive still achieves its best performance at $d=0.1$. For HP C2247 drives, however, the best performance is always achieved by placing the hottest movies in the outermost zone. Part of the reasons is that HP C2247 has a large outermost zone that contains 27% of the cylinders. For HP C2490, the outermost zone contains only about 7% of the cylinders.

In summary, our studies show that the highest performance is achieved by always placing the hottest movies in the two (or three) outermost zones.

4.3. The Effect of Varying Movie Hotness Skew

As described in Section 3, movies stored in video servers have different *temperatures*, and the movie temperature can change over time. In this subsection, we examine the impact of the movie hotness skew on the disk performance under different movie layout strategies, i.e., how much performance benefit we can expect by using the ZBRE layout strategies as movie hotness skew varies.




Figure 4: Favorable Areas of ZBRE.s and ZBRE.1 Strategies.

We first consider the two most promising ZBRE strategies: ZBRE.1 and ZBRE.s. ZBRE.1 tries to fully utilize the high bandwidth of the outermost zone and ZBRE.s trades some bandwidth, which results in an increase of data transfer time, for reduced disk seek time.

Figure 4 shows the favorable areas of the two strategies as functions of data transfer size and movie hotness skew. In general, as the I/O transfer size increases, the data transfer time becomes a more dominant factor, and ZBRE.1 performs better than ZBRE.s. On the other hand, when I/O transfer size is relative small and the movies' hotness is less skewed (the increase of skew degree), the sub-outer-zone strategy, ZBRE.s, becomes more attractive.

The different behaviors of Seagate ST32550 and HP C2490A are mainly due to their physical structures. The Seagate ST32550 is a low-profile drive comprising 6 platters and 3711 cylinders, whereas the HP C2490A is a half-height drive comprising 10 platters and 2582 cylinders. The Seagate ST32550 thus has a higher density and more zones. For the HP C2477 drive, which has the lowest density of the three drives examined (7 platters and 1981 cylinders), the ZBRE.1 strategy is favored.

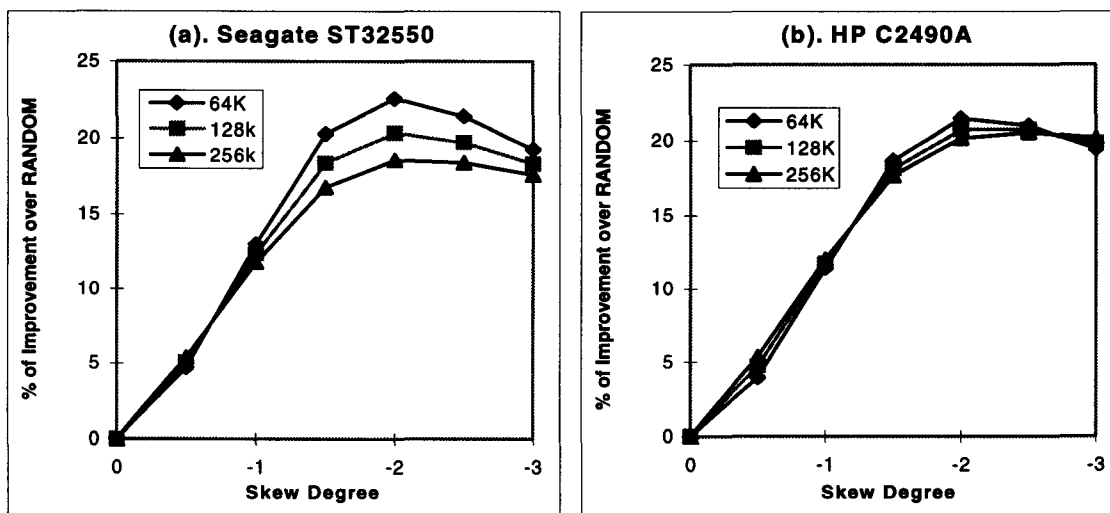


Figure 5: Performance Improvement of ZBRE.1 over the RANDOM Strategy.

Now we compare the performance of our proposed ZBRE family of video layout strategies to the RANDOM strategy under different movie hotness skews. The results are shown in Figure 5 for the Seagate ST32550 and HP C2490A drives using the ZBRE.1 strategy. As shown in Equation (1), when the skew degree $s = 0$, all of the movies have the same temperature. In this case, all of the video layout strategies exhibit the same performance characteristics and a random disk access pattern is observed. When movie hotness is more skewed (the skew degree decreases), the advantages of our ZBRE strategies become more obvious. When 70% to 90% of the viewers select the top 10% of the hot movies, as happens in the real world, ZBRE.1 provides a 17% to 23% per-

formance improvement over the RANDOM strategy, depending on the I/O transfer sizes. In the extremely skewed case, this advantage decreases since, even under the RANDOM strategy, the disk head stays at a few hot movies without incurring many seeks. In this case the advantage comes mainly from the high bandwidth of the outer zones under the ZBRE strategies.

4.4. The ZBRE.r Strategy

In this section, we focus on the ZBRE.r strategy, which trades performance for simplicity of implementation, as described in Section 3. Particularly, we study the performance sacrifice of ZBRE.r for this simplicity as compared to ZBRE.1.




Figure 6: Performance Comparison of ZBRE.1 and ZBRE.r Strategies.

In Figure 6 (a), we plot the mean disk service times for ZBRE.r as compared to the ZBRE.1 and RANDOM strategies. In this figure, the I/O transfer size is 128K bytes and the hot movie skew degree is -2. Figure 6 (b) shows the percentage of improvement over the RANDOM strategy for ZBRE.r and ZBRE.1 as a function of the movie hotness skew degree. From this figure, we see that ZBRE.r has a more decided advantage over RANDOM when the movie hotness is more skewed. In the case of skew degree $s = -2$, ZBRE.r outperforms RANDOM by 16% and ZBRE.1 outperforms ZBRE.r by 5%. On the other hand, when the skew degree is -0.5, ZBRE.r is close to RANDOM, and when the skew degree is -1, it only performs 5% better than RANDOM, whereas ZBRE.1 outperforms RANDOM by more than 12%. Overall, the ZBRE.r performance lies in between RANDOM and ZBRE.1. Its performance is closer to RANDOM when the movie hotness is less skewed, and closer to ZBRE.1 when the movie hotness is highly skewed.

5. Summary

The multiple-zone-recording technique has been used by disk vendors for many years to achieve higher capacities for modern disks. However, it has not been widely used by system designers to improve system performance. In this paper, we presented a family of ZBRE video data layout strategies which lay out movies on disks by taking advantage of the multiple-zone-recording feature. These ZBRE strategies can be widely used in various video or media servers for which some movies or games, etc. are hotter than others. While these strategies were only studied with disks here, they are, nevertheless, directly applicable to RAID devices. The simple ideas developed here can also be extended to traditional file systems or database systems.

We studied the performance of the ZBRE strategies by simulating the behavior of three modern disk drives. Placing hot movies in outer zones and aggregating them together significantly improved the server's disk I/O performance without incurring any extra hardware cost. Our results show that ZBRE layout can achieve an up to 23% improvement of disk service times over a random layout of movies on disks. In addition, within a reasonable movie hotness skew range, ZBRE achieve more than 90% of the outermost zone bandwidth for the three disk drives used in this study.

There are a number of related issues that we plan to investigate in the future. Movie replacement and hierarchical storage management (HSM) are two such examples. Typically the temperature of a movie is a decreasing function of time. Thus, movie replacement strategies are needed when a movie's temperature changes from "hot" to "cold". Similar issues exist in the area of HSM when staging information from one medium to another.

References:

- [1] Yee-Hsiang Chang, David Coggins, Daniel Pitt, David Skellern, Manu Thapar and Chandra Venkattraman, "An Open-Systems Approach to Video on Demand," IEEE Communications Magazine, Vol. 32, No. 5, pp. 68-80, May 1994.
- [2] Patterson, D., G. Gibson, and R. Katz, "A Case for Redundant Arrays of Inexpensive Disks (RAID)," In Proc. of SIGMOD'88, 1988.
- [3] Rangan, P.Venkat and Harrick M. Vin, "Designing File Systems for Digital Video and Audio," In Proc. of 13th ACM Symposium on Operating Systems Principles, (SOSP'91), Operating Systems Review, Vol. 25, No. 5, pp.69-79, 1991.
- [4] Vin, Harrick M. and P.Venkat Rangan, "Designing a Multiuser HDTV Storage Server," IEEE J. on Selected Areas in Comm. Vol. 11, No. 1, pp.153-164, Jan. 1993.
- [5] Rangan, P.Venkat and Harrick M. Vin, "Efficient Storage Techniques for Digital Continuous Multimedia," IEEE Trans. on Knowledge and Data Engineering, Special Issue on Multimedia Information Systems, Aug. 1993.

-
- [6] Keeton, Kimberly and Randy H. Katz, "*The Evaluation of Video Layout Strategies on a High-Bandwidth File Server*," Proc. of the 4th Int'l Workshop on Network and Operating System Support for Digital Audio and Video, Lancaster, England, UK, Nov., 1993.
 - [7] Berson, S., R. Muntz, S. Ghandeharizadeh, and X. Ju, "*Staggered Striping in Multimedia Information Systems*," in Proc. of SIGMOD'94, 1994.
 - [8] Kenchammana-Hosekote, D. and J. Srivastava, "*Scheduling Continuous Media in a Video-on-Demand Server*," in Proc. of the Int'l Conf. on Multimedia Computing and Systems, pp.19-28, May, 1994.
 - [9] Birk, Yitzhak, "*Track-Pairing: a Novel Data Layout for VOD Servers with Multi-Zone-Recording Disks*," submitted for publication. March 1994.
 - [10] Dan, Asit and Dinkar Sitaram, "*Buffer Management Policy for an On-Demand Video Server*," IBM Research Report RC 19347, 1994.
 - [11] Dan, A., D. Sitaram, and P. Shahabuddin, "*Scheduling Policies for an On-Demand Video Server with Batching*," Proc. of ACM Multimedia'94, pp.15-23, San Francisco, Oct. 1994.
 - [12] HP C2490A 3.5-inch SCSI-2 Disk Drives Technical Reference Manual, Ed. 3, Sept. 1993.
 - [13] Zipf, G.K., *Human Behavior and the Principle of Least Effort, An Introduction to Human Ecology*, Addison-Wesley, Reading, Mass, 1949.