Fibre Channel Storage Interface for Video-on-Demand Servers

Shenze Chen, Manu Thapar
Broadband Information Systems Laboratory
HPL-95-125
November, 1995

Fibre Channel,
VOD servers,
I/O performance,
storage interface

The emerging *"information super-highway"* and accelerating improvements in computer and mass storage technologies will soon make multimedia services, such as *video-on-demand* (VOD), a reality. Since a media server will be required to move a large amount of data over distribution networks, its I/O subsystem design is critical to its success. The current most popular I/O interface SCSI has many problems in an I/O intensive environment, because of its limited bandwidth, unfair channel arbitration protocol, etc. In this paper, we explore a new technology, Fibre Channel, as a storage interface for video servers. Specifically, we study Fibre Channel loop topology that attaches multiple disk drives to a server host. In order to compare the performance of the new Fibre Channel interface with parallel SCSI, we built a video server simulator that simulates the behavior of both interfaces. The results show that with the same number of disks and the same system configurations, one Fibre Channel loop can achieve 50% higher performance (measured by the number of concurrent streams supported by the server) than four fast/wide SCSI channels. Our results show that Fibre Channel is attractive as a video/media server I/O interface. We also analyze the buffer size requirements and I/O transfer size for various configurations.

Internal Accession Date Only
To be published in and presented at *Multimedia Computing and Networking 1996 (MMCN '96)*, San Jose, California, January 28-31, 1996.

1. Introduction

The emerging "information super-highway" and accelerating improvements in computer and mass storage technologies will soon make multimedia services, such as video-on-demand (VOD), a reality. To provide such services, servers must continuously read compressed video data from storage and deliver the data to a distribution network in a timely fashion. Similarly, other on-demand services, such as home-shopping, game-on-demand, news-on-demand, etc., can be provided. In contrast to CPU-bound systems, these media servers are actually data movement engines, since their main responsibility is to move a large amount of data from storage to distribution networks. For such servers, the I/O subsystem design is critical to the success of the entire system.

Currently, the most popular disk I/O interface is the Small Computer Systems Interface [1] (SCSI, pronounced "scuzzy"). While SCSI achieves satisfactory performance for many applications such as transaction processing and NFS file servers, it has a few disadvantages when used with media servers. First, SCSI has a limited 20 mega-bytes per second bandwidth. In order to support hundreds of concurrent movie streams that require a large sustained data bandwidth, multiple SCSI channels are necessary. Second, SCSI has an unfair bus arbitration scheme. I/O devices on a SCSI bus have different priorities, depending on their SCSI bus ID. If several devices arbitrate for the bus simultaneously, the one with the highest priority wins. This could cause serious problems for video servers, since in this environment, I/O requests carry time constraints and requests to low priority devices are more likely to miss their deadlines.

Fibre Channel (FC) is a new serial link defined by the ANSI X3T9 Technical Committee as an industry standard [2]. It provides a general transport vehicle for Upper Level Protocols (ULP), such as IPI, SCSI, HIPPI, IP, IEEE 802.2, etc. The bandwidth it provides can be as high as 100 mega-bytes per second, which is five times that of SCSI. Multiple systems or devices can have point-to-point connections through FC switches. In order to avoid the switch costs, Fibre Channel also defines a loop topology that provides a shared media to multiple devices and systems. In an FC loop, each node arbitrates the shared link following a fair arbitration algorithm, as discussed in Section 2.3.

In this paper, we explore the use of Fibre Channel as a storage interface for video servers by running SCSI protocol over the FC loop. The purpose of this paper is to provide a quantitative study on this emerging technology for data intensive applications. In order to compare its performance to parallel SCSI, a video server simulator has been built that simulates both interfaces under the same workload and system configuration. The results show that, while an FC loop provides only a 25% higher bandwidth than four fast/wide SCSI channels, a server with an FC loop can provide a 50% higher performance (measured by the maximum number of concurrent movie streams supported) than its four-SCSI counterpart. This highlights the potential of Fibre Channel replacing SCSI as a future storage interface for I/O intensive systems such as video servers or media servers. By using the video server simulator, we also study some interesting related issues, such as the appropriate number of disks to attach to an FC loop, the buffer memory requirement for each movie stream, the impact of varying the I/O transfer size on system performance, etc. While a transfer size of 256KB was recommended in a previous study [3], the results from our integrated system level video server simulator show that a 256KB transfer size can achieve a slightly better performance than a 128KB transfer size only when the required system buffer memory is doubled. Therefore, we consider a 128 KB transfer size being more appropriate for video servers in terms of the performance/cost trade-off. Lastly, although the initial development cost for FC might be a little high, FC disk vendors have already committed to the price of fast/wide differential SCSI disks when FC drives come to the market.

Both Fibre Channel and video server design are relatively new areas about which little can be found in the literature. Cummings [4] briefly described the general concept of Fibre Channel and the high level architecture of connecting peripherals to systems. Anujan Varma et al. studied the performance of FC switching systems [5,6]. Getchell and Rupert discussed an FC-based local area network [7]. A comparison of FC and 802 MAC services was studied by Ocheltree, Tsai, Montalvo, and Leff [8]. Recently, Anzaloni et al. proposed a design solution for an FC-to-ATM internetworking gateway [9]. For FC loop topology, to our knowledge, no previous work has appeared in the literature. In the video server I/O subsystem area, Rangan et al. [10,11,12] presented a data layout scheme that divides a movie into blocks. These blocks are scattered on disk surfaces, with "gaps" between adjacent blocks. Keeton and Katz [13] presented an interesting scheme to layout video data for playback in different resolutions. This scheme requires video compression algorithms that support decomposition by resolution. Berson et al [14] proposed a staggered method of striping multimedia information across multiple disks based on worst case analysis. Kenchammana-Hosekote and Srivastava [15] studied the issue of scheduling storage devices to guarantee rate requirements for continuous media. This work is also based on worst case analysis. By taking advantage of the multiple-zone-recording feature of modern disk drives, Chen and Thapar [16] studied a way of laying out "hot" movies in outer zones, which achieved an up to 20% of improvement on disk I/O performance. On the same issue, Birk [17] proposed a scheme which pairs each track in an outer zone with a track in an inner zone and places movies on the track pairs. Other work on video server buffer management and multicasting can be found in [18] and [19].

The remainder of the paper is organized as follows: Section 2 describes Fibre Channel and loop topology; Section 3 describes the video server model; the performance evaluation of Fibre Channel interface as well as the SCSI interface is given in Section 4; and Section 5 summarizes this paper.

2. The Fibre Channel Interface

2.1. Why Fibre Channel ?

Because of its rich upper-level protocol support, Fibre Channel can be applied to a wide variety of areas, such as networking, storage interface, system interconnect between clustered workstations, etc. It also allows the direct attachment of mass storage devices to networks, which enables multiple hosts to share storage devices in a distributed environment.

As a storage interface, Fibre Channel has four main advantages over SCSI,

- Fairness:

  The SCSI unfair arbitration problem described in Section 1 may cause serious problems in real-time environments with video servers. However, Fibre Channel, when used as a shared media, does define a fair algorithm for media arbitration. We will describe this algorithm later in Section 2.3.

- High Bandwidth:

  Fibre Channel provides a bandwidth of 100 MB/sec for data transfer, which is five times that of a fast/wide SCSI channel (20 MB/sec).

- Reliability:

  Fibre Channel is a serial interface, which is typically more reliable than a parallel interface (such as SCSI) in terms of signaling, cabling, and physical connectors.

- Distance:

  Fibre Channel allows a distance of up to 10 kilometers between a storage device and its host. This is a critical feature for applications that require duplicate data to be stored in different buildings or even different cities to protect against disasters. In contrast, the distance between a SCSI device and its host must be less than 25 meters.

2.2. Fibre Channel Protocol Stack and Service Classes

Like the ISO/OSI networking protocols, Fibre Channel also defines layered protocols for its data transfer.

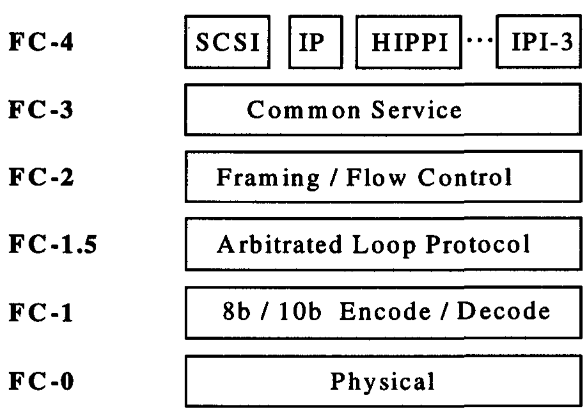

Figure 1: Fibre Channel Protocol Stack.

The protocol stack shown in Figure 1 defines the following layers:

- FC-0 defines the physical layer, which includes the fibre, connectors, and optical/electrical parameters for a variety of data rates. The physical media may be fibre optics, coax, or twisted pair;

- FC-1 defines the transmission protocol, which includes serial encoding and decoding, and error control;
- FC-1.5 defines the Fibre Channel arbitrated loop protocol, which allows the attachment of multiple communication points to a loop without hubs and switches;
- FC-2 defines the signaling protocol, which includes the frame structure, byte sequences, and flow control mechanisms;
- FC-3 defines a set of services, which are common across multiple ports of a node;
- FC-4 defines the mapping between low-level Fibre Channel and high-level protocols such as SCSI , IP, etc.

For example, if a host sends a SCSI command to a storage device (e.g. disk) to read 64K bytes of data, the command block is packetized into a Fibre Channel frame that is sent to the disk. Since the maximum size of a Fibre Channel frame is 2K bytes, the 64KB data is decomposed into 32 frames and sent from the disk to the host in the Data phase. Finally, a Status frame is sent from the disk when the I/O completes.

The current Fibre Channel standard defines three classes of service, which are distinguished primarily by the methodology used to allocate and retain the communication circuit between nodes.

- Class 1 -- Dedicated Connection:

  Dedicated connections must be established before transferring data frames. Once a connection is established, the bandwidth is guaranteed until one party releases the connection.

- Class 2 -- Multiplex:

  This is a connectionless service on a frame-by-frame basis. The Fabric, if present, routes the frame to the destination node. An end-to-end acknowledgment is required in this class of service.

- Class 3 -- Datagram:

  This is also a connectionless service and is distinguished from the Class 2 service by not requiring the end-to-end acknowledgment.

In addition, Fibre Channel also defines an *"intermix"* of services that interleave Class 2 and Class 3 frames during an established Class 1 connection. In this case, Class 1 frames have a higher priority than Class 2 and Class 3 frames.

2.3. Fibre Channel Loop Topology

Since the Fibre Channel arbitrated loop provides a cost effective way of connecting multiple disks to video servers without switches, in this subsection, we briefly describe the loop topology and the loop arbitration algorithm.

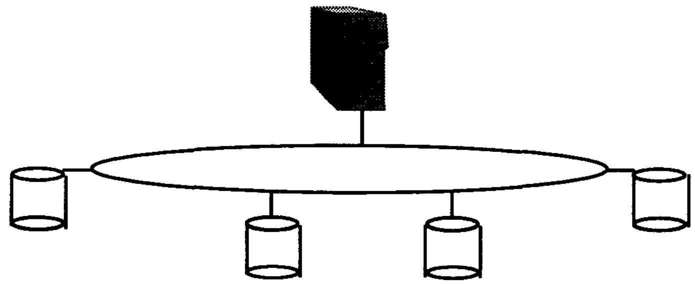

Figure 2: Fibre Channel Arbitrated Loop Topology.

Figure 2 shows a Fibre Channel Loop with five nodes, one host and four disks (multiple hosts are allowed). Each node has a port consisting of a transmitter and a receiver. The loop topology follows the same Fibre Channel point-to-point communication philosophy, i.e., only two nodes in the loop can communicate with each other at a time. Since all of the nodes share the same link, an arbitration is required before a node is allowed to hold the link for transmission.

Unlike the SCSI priority arbitration rule, Fibre Channel defines a fair arbitration algorithm, which includes an access window. Whenever a node wants to arbitrate, it joins the window. If a node wins arbitration once, it must wait until all the other nodes in the current window have had a chance to win an arbitration. At that time, a new window is started, and each node has a chance to arbitrate again. This guarantees that each disk has an equal chance to transfer data. The only exception is that hosts can be assigned a high priority during the loop initialization phase. This allows hosts to send commands to multiple disks promptly.

3. The Video Server Model

In order to study the behavior of the Fibre Channel interface, we built a video server simulator that simulates the activities of a server in detail. The video server model is shown in Figure 3.

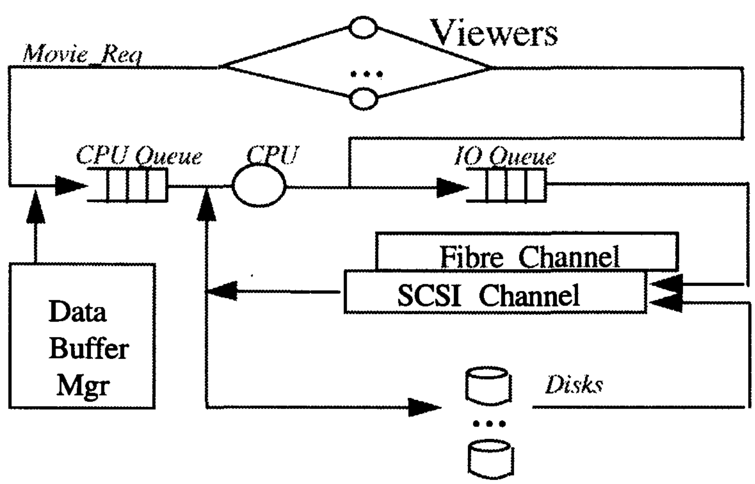

Figure 3: The Video Server Model.

3.1. Movie Data Layout

In this model, movies are striped across all the disks attached to the server, with a striping depth equal to the I/O transfer size. For example, in Figure 4 a movie is decomposed into segments, *A1, A2,...An,* with each segment stored on a separate disk (the segment *Ak+1* is wrapped around to Disk 1 again). During play back, the server issues the first I/O request to Disk 1, the second request to Disk 2, ..., at certain time intervals determined by the stream bit-rate and the buffer size allocated to the stream as discussed in Section 3.2. Each movie is pre-allocated to contiguous cylinders at certain location on the disks according to a layout strategy [16]. All I/O requests generated by a stream are directed to its cylinder locations.

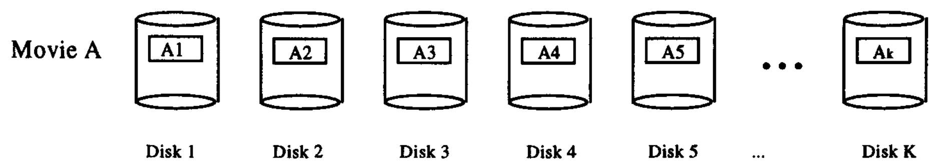

Figure 4: Data Striping across all Disks.

3.2. Buffer Management

When a viewer movie request arrives, the server allocates a main memory buffer to that stream. The buffer space is divided into segments that have the same size as those described in Section 3.1. Movie segments are pre-fetched and played back in round-robin order. When a movie segment has been played-back, the server issues the next I/O request immediately to fill the free buffer. Thus, with a constant bit-rate (which is assumed in the simulation), the generation of I/O requests for each stream is triggered by periodic signals, and each I/O request carries a deadline which is equal to the total play-back time of all the remaining segments assigned to the stream. Obviously, the play-back time of each segment is a function of the stream bit-rate and the segment size. At a 3 Mb/sec bit-rate and 128KB segment size, if the stream is assigned a 512 KB buffer space, then I/O requests are issued at periods of 333 milliseconds and each I/O request for this stream carries a deadline of one second.

3.3. System Workload and I/O Access Patterns

In this model, multiple viewers randomly generate movie requests. Each request selects a particular movie according to a movie hotness distribution. Movie lengths are uniformly distributed between 100 and 120 minutes. A movie request can also select a play-back rate for different qualities of display. For simplicity, however, the results reported below use a fixed stream bit-rate of 3 Mbit/sec, and there are no fast-forward or fast-backward operations. When a movie finishes, the viewer returns to a *"think"* state for a random period of time and then generates another movie request.

By using the above data layout and buffer management strategies, for each disk, say Disk 1, a conceptual service round consists of the period between the arrival of request *A1* and request

Ak+1. All requests arriving at Disk 1 during this period belong to different streams. Therefore, disk seeks are expected between servicing requests. This results in a "*large random access*" pattern on each disk, where "large" is relative to the accesses to traditional file servers and database systems. Since at any time, only a fraction of the total streams may have I/O outstanding to Disk 1, and the *SCAN* algorithm is used for disk queues, this results in a scheme similar to the "Group Sweeping Scheme" (GSS) [20], but with moving group boundaries.

3.4. Storage Interface

For the storage interface, which is the most interesting part of this study, we simulated the Fibre Channel SCSI protocol FCP in detail. We assumed that one Fibre Channel loop was attached to the server with 1G bit per second of bandwidth, multiple disks were connected to the loop, and the loop length was less than 20 meters. The simulator simulated the Fibre Channel loop arbitration, loop open and close, as well as the SCSI Command, Data, and Status phases carried over the Fibre Channel link. Class 3 service was adopted. We also simulated four parallel fast/wide SCSI channels for comparison. Movies were striped across the four SCSI channels so that each channel received the same workload. Since many server platforms do not support an odd number of SCSI channels, we used four SCSI channels instead of five for practical reasons.

4. Performance Results

In this section, we report the results obtained from the video server simulator described in Section 3 above. The performance metric used was the total number of concurrent streams the server could support. We varied the number of disks attached to the server and the total amount of buffer space assigned to each stream. The results were compared to systems using four fast/wide SCSI channels with the same number of disks and the same size of buffers.

Table 1: Simulation Parameters.

System Parameters:	
Movie Stream Bit Rate	3 Mbit /sec
Movie Length	100 - 120 min.
Num. of FC Loops	1
Num. of F/W SCSI Channels	4
Num. of Disks	4 - 60
I/O Transfer Size	64K, 128K, or 256K
Buffer Size per Stream	256K - 2M
Disk Parameters:	
Avg. Seek Time (ms)	10.5 ms
Data Transfer Rate	3.1 - 5.3 Mbytes /sec
RPM	5400

For simplicity, we assumed that each movie had the same probability of being selected. Table 1 lists several of the simulation parameters. The disk parameters are from current typical drives. We also simulated some of the state-of-the-art disk drives with different parameters. Their results are reported in Section 4.2.

The performance results identify the maximum number of concurrent movie streams that can be serviced without any I/O request missing its deadline. Each data point is obtained by averaging 10 runs. Within each run, 500 movies are viewed by multiple viewers. Since a movie lasts 110 minutes on the average, at a 3 Mbit/sec stream rate and a 128KB I/O transfer size, each movie stream generates about 20,000 I/O requests. Hence, each data point obtained from the simulator corresponds to 10^8 I/O requests issued without a single loss.

4.1. Fibre Channel vs. SCSI

Figure 5 illustrates the performance results of comparing one Fibre Channel loop with four fast/wide SCSI channels. Figure 5(a) shows a 512K buffer allocated to each stream, and (b) a 1MB buffer allocated to each stream, both with an I/O transfer or segment size of 128KB.

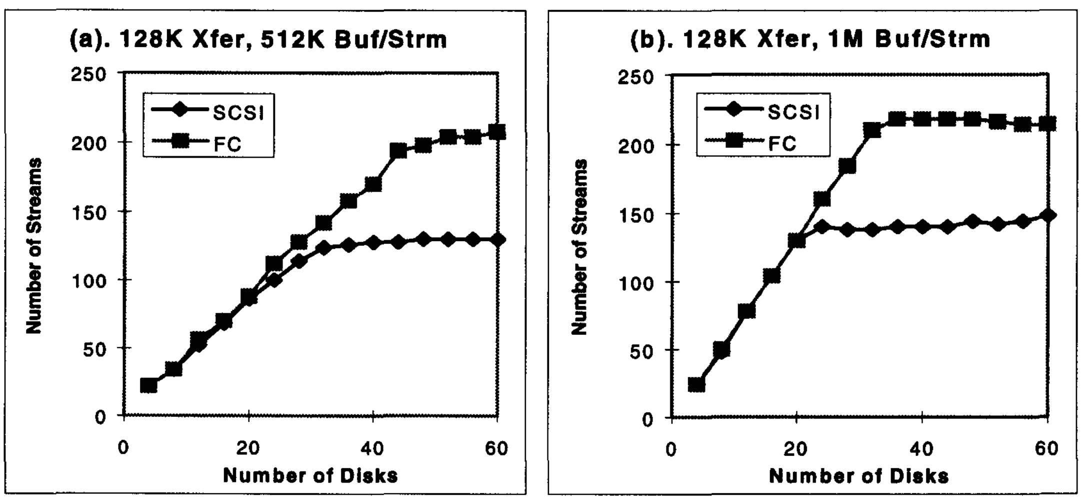

Figure 5: Performance of Fibre Channel vs. SCSI.

From the figure, we see that when the number of disks is less than 20, there is no difference between one Fibre Channel loop and four fast/wide SCSI channels, since the disks are the bottleneck when channel utilization is relatively low. When more disks are added, channel contention increases, and the SCSI low bandwidth and unfair arbitration problem become more dominant, causing the SCSI channel to become a bottleneck. Adding more disks increases storage capacity, but does not help system performance. On the other hand, the Fibre Channel link becomes saturated when 40 disks are connected to the loop. Overall, while the FC loop provides a 25% higher bandwidth than four fast/wide SCSI channels in this configuration, the FC loop could perform 50% to 60% better than four fast/wide SCSI channels. Out of the 100MB theoretical peak bandwidth, an FC loop can actually achieve 80MB/sec of data bandwidth, and the remaining bandwidth is taken by various overheads, such as arbitration, open, close, etc. On the other hand, the four fast/wide SCSI channels can achieve 55.5MB/sec of useful bandwidth (out of 80MB/sec theoretical peak). Therefore, for five SCSI channels, which provide the same theoretical peak bandwidth as one FC loop, we predict the maximum achievable data bandwidth to be 69.4MB/sec. The difference of 10MB/sec between one FC loop and five SCSI's is mainly caused by the FC fair arbitration scheme.

4.2. Using Fast Disk Drives

In this section, we explore the performance impact of using higher performance disk drives. The disk drive parameters [21] are listed in Table 2 and the results for the fast drives are plotted in Figure 6. The fast drives were used for both the Fibre Channel and the fast/wide SCSI's.

Table 2: Disk Parameters.

	Typical Drive (D1)	**Fast Drive** (D2)
Avg. Seek Time (ms)	10.5	8.34
Data Transfer Rate	3.1 - 5.3	5.9 - 9
RPM	5400	7200

From the figure, we observe that the peak performance for the Fibre Channel loop remains the same for both the faster and the slower drives. The number of disks needed to achieve peak performance, however, is reduced for the faster drives. While more than 30 typical disks are needed to reach peak performance, 20 fast disks may saturate the Fibre Channel link. Again, increasing the number of disks provides more storage capacity, but does not help system performance. The figure also shows that adding more disks to the loop hurts system performance since there is a fixed per-node delay for all information transferred through the Fibre Channel loop. Based on the results presented here, we can expect that in the next few years, as disk performance increases further, even fewer disks (10 or less) will saturate a Fibre Channel loop.

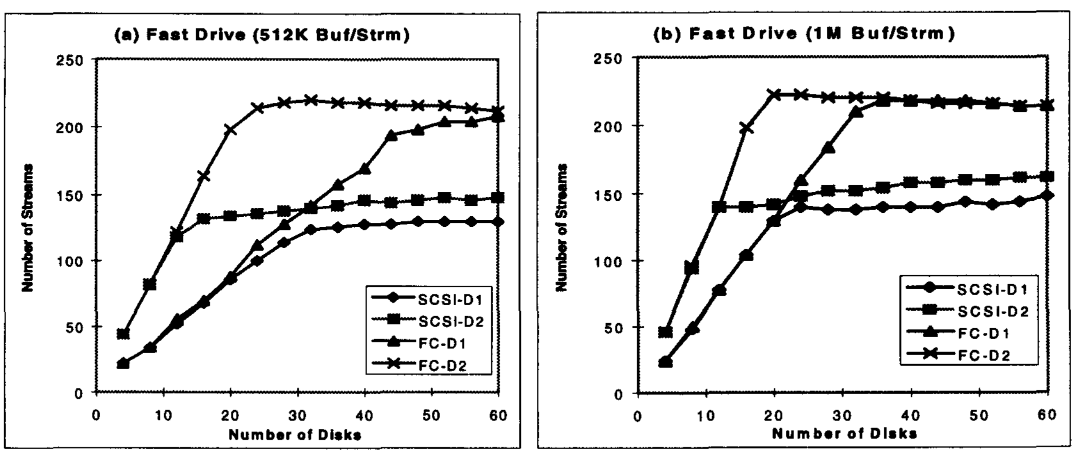

Figure 6: Using Fast Disk Drives.

4.3. Varying Buffer Sizes

In the previous figures, we fixed the buffer size at 512K bytes or 1M bytes per stream. Now we study the performance impact of varying buffer sizes for each video stream. The results are shown in Figure 7, with one Fibre Channel loop, four SCSI channels, and a 128KB transfer size. The number of disks used was fixed at 16 and 32 for both Fibre Channel and SCSI systems. As can be seen from the figure, there is virtually no performance difference between Fibre Channel and SCSI when 16 disks are used. In this case, the disks are saturated when the channel utilization is low. The optimum buffer size required is only 768K per stream. When we increase the number of disks to 32, a significant performance difference between Fibre Channel and SCSI is observed. In terms of buffer requirements, the Fibre Channel system needs 1M per stream to achieve peak performance, while the SCSI system reaches its peak by using 512K to 768K buffer per stream. Using more buffer space increases the system cost without any performance improvement and is not recommended.

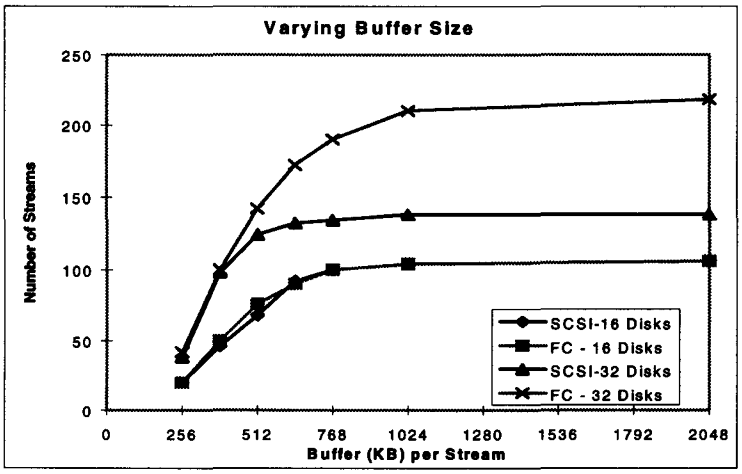

Figure 7: The Performance Impact of Using Different Buffer Sizes.

4.4. Different I/O Transfer Sizes

What is the proper I/O transfer size for video streams? Intuitively, a larger transfer size per I/O will increase the channel as well disk utilization, but may require more buffer space. A smaller transfer size generates more control overhead, such as command processing, multiple interrupts, and possibly more disk seeks. The benefit of using a smaller transfer size is a reduction in the buffer space required for each stream. For SCSI systems, a smaller transfer size also helps to alleviate the SCSI bus arbitration unfairness problem. In Figure 8, we plot the number of concurrent streams as a function of buffer size under different I/O transfer sizes. The number of disks is fixed at 32 typical drives as mentioned above. For Fibre Channel system, a relatively small transfer size of 64K performs best when the buffer size per stream is less than 640K. When the buffer size per stream is greater than 640K and less than 2M, the 128K transfer size is the best. The 256K transfer size becomes best only when there is a large buffer available that can allocate 2M buffer for each video stream. This implies that, for a server capable of supporting 200 concurrent video streams, 400M bytes or more of buffer space is required. A similar situation is also observed for the SCSI interface.

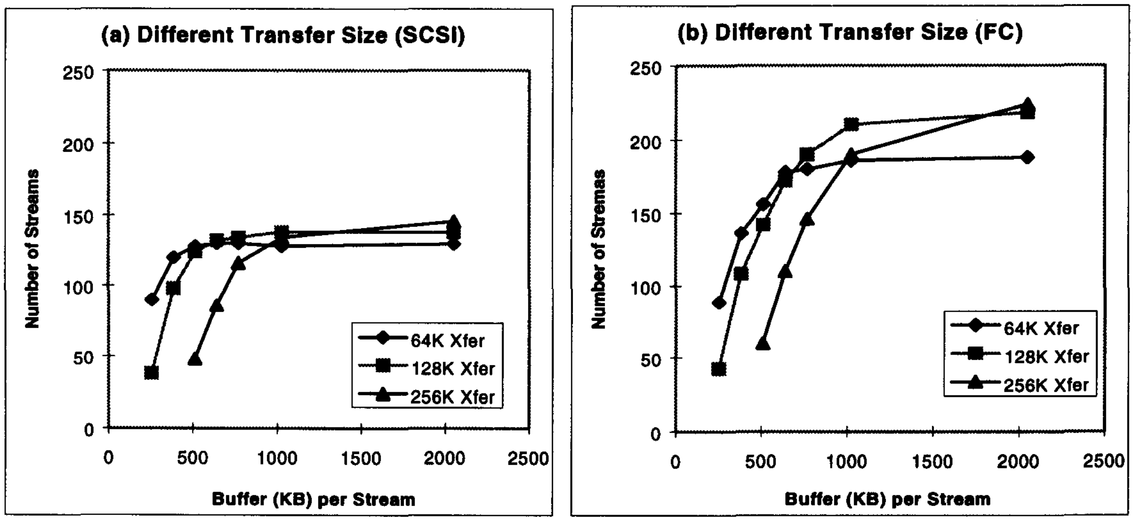

Figure 8: Using Different I/O Transfer Sizes.

5. Summary

In this paper, we studied the Fibre Channel loop as a mass storage interface for video servers that require real-time disk I/O. We chose the Fibre Channel loop topology instead of switches in order to avoid the high cost of a Fibre Channel fabric. Fibre Channel, a newly defined industry standard, has many advantages over SCSI, currently the most popular storage interface. First, Fibre Channel provides 1 Gbit/sec of bandwidth, which is five times that of fast/wide SCSI. Secondly, Fibre Channel loop defines a fair arbitration algorithm for the shared link, while SCSI suffers from its priority-based arbitration protocol, which causes I/O requests directed to low priority devices to be more likely to miss their deadlines in a real-time video server environment. In addition, a serial Fibre Channel link is more reliable than a parallel SCSI channel, and Fibre Channel allows a distance of up to 10 kilometers between storage devices and hosts, as compared to the maximum of 25 meters allowed by SCSI.

The performance of Fibre Channel storage interface has been studied by using a video server simulation model, which includes many components such as memory management, CPU scheduling, disk scheduling, storage interface, etc. For the Fibre Channel interface, we simulated the SCSI protocol FCP running over FC links with a Class 3 service. Specifically, we simulated loop arbitration, open and close, as well as each of the SCSI Command, Data, and Status phases. In order to compare Fibre Channel with the SCSI interface, we also simulated four fast/wide SCSI channels with the same workload and system configurations. The results showed that the maximum number of concurrent movie streams supported by one Fibre Channel loop could be

50% higher than that delivered by four fast/wide SCSI channels. This makes Fibre Channel very attractive as an I/O interface for video or media servers.

Finally, we have developed several video data layout strategies and real-time disk scheduling algorithms, and have integrated them into the above video server simulator. Our next step is to test these strategies and algorithms in an integrated environment. While a zero I/O loss ratio was targeted in this study, we also plan to relax this requirement and study the system performance impact under certain acceptable loss ratios as part of our future work.

References:

[1] ANSI Standard X3T9.2/86-109 *Small Computer System Interface-2 (SCSI-2),* Oct. 1991.

[2] ANSI Standard X3T9.3 *Fibre Channel Physical and Signaling Interface (FC-PH),* Rev 4.0, May 1993.

[3] Reddy, A.L. Narasimha and J. Wyllie, *"I/O Issues in a Multimedia System,"* COMPUTER, 27(3):69-74,1994.

[4] Cummings, Roger, *"System Architectures Using Fibre Channel,"* Proc. 12th IEEE Symposium on Mass Storage Systems, pp.251-256, 1993.

[5] Varma, Anujan, Vikram Sahi, and Robert Bryant, *"Performance Evaluation of a High-Speed Switching System Based on the Fibre Channel Standard,"* Proc. of the 2nd IEEE Int'l Symposium on High-Performance Distributed Computing, Spokane, Washington, July 1993.

[6] Varma, Anujan, Shree Murthy, and Robert Bryant, *"Using Camp-on to Improve the Performance of a Fibre Channel Switch,"* Dept. of Computer Engineering, Univ. of California/Santa Cruz, CA 95064, 1993.

[7] Getchell, D. and P. Rupert, *"Fibre Channel in the Local Area Network,"* IEEE LTS, Vol.3, No. 2, pp.38-42, May 1992.

[8] Ocheltree, K.B., T.C. Tsai, R. Montaivo, and A. Leff, *"A Comparison of Fibre Channel and 802 MAC Services,"* Proc. of IEEE , pp.238-246, 1993.

[9] Anzaloni, A., M. De Sanctis, F. Avaltroni, G. Rulli, L. Proietti, G. Lombardi, *"Fibre Channel (FCS) / ATM Internetworking: A Design Solution,"* Proc. Global Telecommunications Conf., Vol.2, pp.1127-1133, 1993.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] Berson, S., R. Muntz, S. Ghandeharizadeh, and X. Ju, *"Staggered Striping in Multimedia Information Systems,"* in Proc. of SIGMOD'94, pp.1-12, 1994.

[15] 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.

[16] Chen, Shenze and Manu Thapar, *"Zone-Bit-Recording-Enhanced Video Data Layout Strategies,"* Proc. of MASCOTS'96, San Jose, Jan. 1996.

[17] Birk, Yitzhak, *"Track-Pairing: a Novel Data Layout for VOD Servers with Multi-Zone-Recording Disks,"* Proc. of Int'l Conf. on Multimedia Computing and Systems, Washington DC, May 1995.

[18] Dan, Asit and Dinkar Sitaram, *"Buffer Management Policy for an On-Demand Video Server,"* IBM Research Report RC 19347, 1994.

[19] Dan, Asit, Dinkar Sitaram, and Perwez Shahabuddin, *"Scheduling Policies for an On-Demand Video Server with Batching,"* IBM Research Report, RC 19381, 1994.

[20] Yu, P.S., M.S. Chen, and D.D. Kandlur, *"Grouped Sweeping Scheduling for DASD-based multimedia storage management,"* Multimedia System Journal, 1:99-109, 1993.

[21] Seagate Product Overview, March 1994.