CMPSCI 754 Multimedia Systems **Spring 2001**

Lecture 5: Feb 15

Lecturer: Prashant Shenoy *Scribe: Zihui Ge*

5.1 Multimedia Storage Servers

5.1.1 Overview

Different from normal text based storage system, multimedia storage servers are 1) larger in capacity; 2) oriented for handling large file (for example, 1 hour mpeg1 video file needs $\sim$ 1GB); 3) required to provide real-time service instead of best-effort service for text.

The fundamental issues for multimedia storage server design include structures for multimedia file placement(on storage device), fault tolerance against disk failure, disk scheduling algorithms, retrieval and admission control to accomplish realtime application's specification, and I/O stream sharing scheme (such as buffering, batching, caching problems).

5.1.2 Disk fundamentals

For each disk operation, the response time is composed of three components:

$$\text{Response time} = \text{Seek time} + \text{Rotational latency} + \text{Transfer rate}$$

- Seek time: the time for the head to move from current track to the track wanted. For modern disk, the seek time is less than 5ms.
- Rotational latency: the time for the sector wanted rotate to the position under disk head. For example, the average rotational latency is 5ms for a 6000RPM disk.
- Transfer rate: time to access the data wanted. Disk tranfer rates are about 7-10Mb/sec.

Some scheduling polices for single disk controller:

- FCFS (First Come First Serve) scheduling: advantage is no starvation, but is less optimized in average response time.
- SCAN: disk head moves back and forward between the innermost track and the outermost track and picks up the requests on the track it passes.
- SSTF (Shortest Seek Time First) scheduling: it may lead to starvation for requests on intermost or outermost tracks.
- SATF (Shortest Access Time First) scheduling: similar to SSTF, but also consider rotational latency and transfer time.

5.1.3 Disk array fundamentals

Disk array system is composed of a collection of disks(4-12 in small systems, serveral hundreds in large system) and a disk array controller which presents the whole array of disks as a single disk to the operating system(transparency). Using disk array not only increases total storage capacity but also improves system throughput.

Striping is the technique of interleaving the storage of each multimedia stream among disk. Instead of storage the whole file on a single disk, it breaks the file into blocks and places the blocks on different disks.

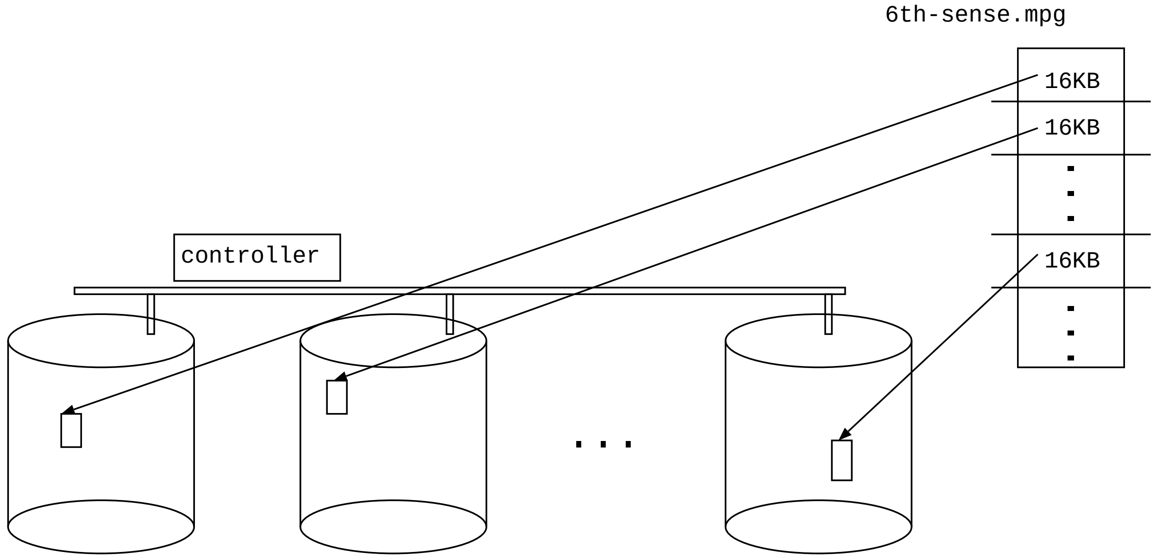

Figure 5.1: Striping

Stripe unit is the maximum amount of logically contiguous data that is stored on a single disk. Degree of striping is the number of disks across which a media stream is striped (not necessarily be the total number of disks). The advantages of using striping are providing the possibility of parellel access, improving throughput and load balancing.

Fault tolerance consideration is important in disk array system in that the MTBF(Mean Time Between Failure) of the disk array is approximately the MTBF of the disk divided by the number of disks in the system. For example, in a disk array system of 200 disks, if the average MTBF per disk is 500,000 hours, the MTBF of the system is only 2500 hours(14 weeks).

5.1.4 Video Storage Server

Two possible server architectures:

- Client-Pull
  Server retrieves data only in response to explicit requests from client. In order to ensure realtime playback continuity, client should determine the playback instant of a frame, estimate response time for each request and issue request accordingly. However, this is insufficient because the response time is very sensitive to the system load, of which the client has no knowledge.

- Server-Push
  Client provides the playback rate information and server periodically "generate" disk operation requests for each client. In each round, a fixed number of frames is retrieved for each media stream. For real-time playback requirement, server should guarantee that the total service time will not exceed the playback duration of frames retrieved during a round.

To design a video server, (e.g. determine the stripe unit size, degree of striping, disk scheduling policy, etc.), one should count in factors such as the number of disks in the disk array, the number of clients accessing the server... The objective is to maximize the number of clients that can be simultaneously supported.