Homework Turned In

Turned in cover letters, resumes, thank-you notes.

Homework Assignment

Two proofs of subscriptions. Worth 20 points. Ask about this since Steve may have discussed this before class and I got to class just before it started. He forgot. 2 printed magazines or 5 email subscriptions.

Video Server

To serve means: to be a slave, to do the bidding of others, providing a service, impartially doing what is right.

First video server introduced around 1994. Digitized some video and put it on a server.

Isochronous

Eye-SOC-hron-ous is pronounciation. Signals that are periodic signals in which the time interval that separates any two occurrances is equal to some unit interval. An example of this is sync in video. Digitized voice signal samples occur isochronously at precisely the sampling interval.

Asynchronous no common clock. Synchronous is one common clock. Isochronous can have multiple clocks (multiple data rates) on the same line at the same time.

Packet data signals are not isochronous.

Isochronous is not as rigid as synchronous service, but not as lenient as asynchronous service.

What is a Video Server?

The first is **video capture**. You must have a way to record it onto the server. As content is created digitally, it may be that capture becomes a totally separate function. Loading video onto servers may be through file transfers.

The next component to a video server is **playback**. What a true video server, the output is one or more traditional video streams.

If you can capture and playback video streams, there has to be a **structure to the information**, the method used to identify, store and request video information. This may be a clip's name, or a database with searching capability to identify clips of interest. Camera person had to be careful to take a good establishing shot as the first take in a segment, since it would be reflective of the name of the clip.

How long to transmit from truck to station, do some edits, and put clip on server? Stations wanted this to happen in a matter of minutes. Each person doing their separate jobs could be able to do them at the same time. That is the point of the server. Anybody could get access to a particular clip and do their work on it. Anybody could look at it and see how it looks. Can be broadcasting content out while it is being captured, conceivably. Camera person had to think in terms of what should be the sequence in the final cut, and take the shots in that order, in order that the shots can be transmitted and edited on the fly in the order needed.

A video server cannot take any time to do something; it must happen immediately. Unlike a data server, which one can expect to take a while to perform the data processing function requested. There is no time to wait for things in the video world, unlike the data world. Can't be dropping frames on a video system.

Finally, you can't have a video server without **connectivity**, a video server must be able to *broadcast* its information to multiple locations.

Difference between master control and TD. TD is production control, for a particular production (tape, live, server, etc). Master control can switch between production control output and recorded content like commercials. Master control does the final control of what goes out of a TV station.

A video server is not merely a high capacity hard drive that stores moving video and other associated data but rather a highly specialized device. Other data includes: audio, time code, VITC, metadata that travels with the production. Metadata like who made it, how long it is. Also called ancillary data.

The video server is the multimedia equivalent of a data file server, having a full computer system, including storage, processing, and input/output functions, and video and audio encoders and decoders to convert analog content into digital and vice versa.

There is no horizontal or vertical intervals in digital. So, there's no place for VITC or closed captioning like there used to be. They are stored elsewhere.

We use SDI (serial digital interface) in the studio.

IT environments are virtually asynchronous and video requiresisochronous behaviors, maintaining sufficiently high throughputs – so all channels can play-out or record continuously – requires sophisticated designs and complex hardware configurations.

The parallel synchronous movement of large blocks of data is needed for video – and this is what typically sets video server architectures apart from IT-based servers and storage systems. Data uses 7 or 8 bits, video based on 10 bit system (per channel). 2 to the 8^{th} power is 256. 2 to the 10^{th} power is 1024.

Large number doesn't tell definition as much as does the sample rate. What is the 1024 bits for? Luminance levels in active video, running from setup (7.5 IRE) to peak white (110 IRE). Represent setup as 0000000000 and peak white as 1111111111, for example. 110 IRE gives us a cushion so that if somebody makes a mistake and gets white higher than 100 IRE, there's still room for error. Very finely measuring at each sample point, and generate a 10-digit number at each sample point. Headroom necessary to avoid rollover from excessive white to numbers representing black (oneplus 1111111111 is 0000000000). Headroom analogous to the -18 dB or -20 dB recording level in digital audio.

If just used 8 bits, you might notice the steps between levels.

Video is represented using 10 bits. Remember this and put in black book.

Sample Rate

Sample rate is how often we generate the 10 bits during the video signal. Will cover this in a later lecture. We will cover it a little bit here at Ralph's request. In NTSC, burst is 3579545 Hz. Nice to make our sample frequency a multiple of something we carry with video (like burst). Nyquist limit says that you have to sample at a rate at least 2 times the highest frequency being sampled. Highest frequency in video is 4.2 MHz (the 3 dB rolloff at the high end). So, sample rate is 8.4 MHz. What can you multiply 3579545 by in order to get 8400000? About 2 1/3 will do it, but not so easy to do. 3 times would be OK. In early days, decided to go 4 times subcarrier just to be safe/sure (**14.318181 MHz**). Remember this number; don't round it off. Crystals are cut for this frequency. Called **four times subcarrier**, also written as **4FSC** or **4F**$_{SC}$. This is locked to burst if you arereceiving the signal.

Sampling 10 bits every cycle of 14.318181 MHz. For color, since we don't see it as well as luminance (I and Q are substantially less than half of our luminance video). We sample the entire composite video stream at 14.3181818 MHz, get 10 bits each time. What do we do with the 10 bits? This includes the luminance and chrominance. This is digital composite, not component. 69.8 nanoseconds is the period at this frequency. How to deal with 10 bits every 69.8 nanoseconds. 143 million bits every second. Studio doesn't compress the video until it gets to the transmitter.

SDI (serial digital interface). Convert the 10 bits of parallel data into serial data. Use a shift register clocked at 10 times that rate, or 143 megabits per second (143 MHz).

What sort of cables used to transport this data around the station? 75 ohm coax cable. Not ideal, but gets the job done as long as we don't have to use long runs. The D-2 recorders run at 143 megabits per seconds. Encode everything: burst, sync tip, setup. So, when decode it, haven't messed with anything. This is composite digitizing.

This is fast being replaced by component digital video.

At the time, 143 MHz was hard to achieve. That's probably why they accepteda limit of 4 times color burst as the frequency.

Instead of locking right onto burst as is, they shift the burst phase into what Q would be and then start the digitizing.

Back to Video Servers

Video servers can range from a standard PC for small-scale systems to massively parallel devices.

What's inside of a Video Server?

RAID (redundant array of intelligent/inexpensive drives/disks). Why do this?

Access Time

If we were talking about a single disk, access time would be the time the disk took to receive the command, move the heads to the correct position, wait for the disk to rotate to the correct location and start to read.

Access time is particularly important to multi-user and database environments. These environments have many requests to the same data.

Access time is important when deciding what level of RAID you will implement.

Array

An array is a grouping of disk drives.

There could be two disk drives in the array or there could be 100 disk drives in the array.

In our server, there are 4 drives that store the video data. Need for high availability (read while writing).

Cost per Megabyte Penalty

Term used to indicate what the percentage of total storage is dedicated to providing data redundancy.

Chunk

Breaking up of data so that it can be stored across multiple disk drives.

Head oscillates to some degree because it can't stop exactly, being mechanical. This is settle time. While one finds the track and settles, it is ready to record.

Breaking up into chunks done by controller. Imagine breaking up a 80 KB file into five 16 KB pieces. These 16 KB pieces would be referred to as **chunks**.

Another commonly used term is **stripe**. Stripe is the preferable term.

Controller part of the RAID. Can even detect disk failure.

Disk Set

Specific number of drives grouped together with a single characteristic (e.g., RAID 0, RAID 5).

In the MPEG server in the back, there are two RAIDs. One is for the operating system and the data part of the server, the other is for the video part of the server.

A disk set can encompass a whole array or be a subset of the array.

There can be multiple disk sets in an array. A disk set will present itself to an operating system as an individual disk drive.

Disk set is also referred to as **RAID set**.

Duplexing

When members of a disk set are spread across different SCSI buses. Small computer systems interface. Related to computer bus. Duplexing is for speed.

This is important for two reasons:

1. it relieves dependence on one SCSI bus in the event of failure
2. it increases performance by moving data across two different buses simultaneously

If the controller is a smart controller, it will recover for you. Supports hot-swapping (pulling out a live drive while it is running and dropping in a new drive). System will continue to run with a bad drive, less optimally, but will keep you on the air.

Mirroring

Data is stored twice (or more) on two or more different disk drives. If one drive fails, can still use the other drive. Important for continuity and staying on the air.

Shadowing is a commonly shared term with mirroring. What is written on one disk is written on another disk.

Parity

Mathematical equation that allows data to be checked for integrity. It is a check number, calculated based on the data. Can compare the data by recalculating the checknumber later and seeing if it and the original check number matches.

The last number on your credit card number is a parity; that's how they checkto see if you entered the correct number.

Simple example. Parity bit is zero if binary word has an even number of ones and one if binary word has an odd number of ones. For example, 101101_2 would have a parity bit of zero.

Data is generated that allows the stored data to be checked for integrity.

Parity Disk

RAID levels 3, 4, and 5 have a dedicated drive for storing parity infomation. Parity infomation may stay on one disk (3, 4, 5) or move between disks as parity chunks.

Partition

Breaking up of a disk set into smaller segments. The smaller segments will appear as individual disk drives to the host, while still maintaining the RAID properties of the disk set.

Striping

RAID accomplishes the formation of a single logical drive from multiple units through a process called **striping**.

Striping involves logically arranging infomation so that individual files are spread among a group of drives. A stripe segment can be as small as a single byte, or as large as multiple sectors.

Striping has the advantage of faster data access because individual drives can be accessed in parallel (the data on the computer data bus is in parallel). The disadvantage is that the array capacity, once formatted, is fixed, meaning expansion is not as simple as adding drive modules.

Striping a video stream across *multiple* servers is commonly called **wide striping**.

Transfer Rate

How fast the RAID set or subsystem can transfer the data to the host. Important when large contiguous blocks of data are being used.

Video and image files are examples of large contiguous file transfers that occur in streaming mode.

Transfer rate is important when deciding what level of RAID you will implement.

Mirroring isn't as good for video since writing the same data to multiple drives slows it down.

RAID 0

Writes data across the drives in the array,one segment at a time.

Striping offers high I/O rates since read/write operations may be performed simultaneously on multiple drives. For example, disk 1 takes blocks 1, 4, and 7, disk 2 takes blocks 2, 5, and 8, and disk 3 takes blocks 3, 6, and 9.

Term	***Definition***
Random Read Performance	Very good; better if using larger stripe sizes if the controller supports independent reads to different disks in the array
Random Write Performance	Very good, again, best if using a larger stripe size and a controller supporting independent writes
Sequential Read Performance	very good to excellent
Sequential Write Performance	very good
Cost	lowest of all RAID levels
Special Considerations	Using RAID 0 array without backing up any changes made to its data at least daily is a loud statement that that data is not important to you
Recommended Uses	Non-critical data (or data that changes infrequently and is backed up regularly) requiring high-speed, particularly write speed, and low cost of implementation. Audio and video streaming and editing; web servers; graphic design, high-end gaming.

RAID 1

Writes data to two drives simultaneously.

If one drive fails, data can still be retrieved from the other member of the RAID set.

This process is also called mirroring. The most expensive RAID option but offers ultimate in reliability.

Term	*Definition*
Random Read Performance	
Random Write Performance	
Sequential Read Performance	
Sequential Write Performance	
Cost	lowest of all RAID levels
Special Considerations	Using RAID 0 array without backing up any changes made to its data at least daily is a loud statement that that data is not important to you
Recommended Uses	Non-critical data (or data that changes infrequently and is backed up regularly) requiring high-speed, particularly write speed, and low cost of implementation. Audio and video streaming and editing; web servers; graphic design, high-end gaming.

RAID 0+1

A combination of striping and mirroring. Provides optimal speed and reliability but possesses the same cost problem as RAID 1.

RAID 3

Parity on separate disk.

The BTS Media Pool uses RAID 3.

RAID 5

Employs combination of striping and parity checking. Use of parity checking provides redundancy without double disk capacity.

Parity checking involves determining whether each given block has an odd or even value. These values are summed across the stripe sets to obtain a parity value. With this parity value, the contents of a failed disk can easily be determined and rebuilt on a spare drive.

Parity across disks.

Multiple/combined RAID Levels

A multiple RAID level is generally created by taking a number of disks and dividing them into sets. Then, the second RAID level is applied to the arrays to create a higher-level array. This is why these are sometimes called *nested* arrays.

Since there are two levels, there are two ways they can be combined. The choice of which level is applied first and which second has an impact on some important array characteristics.

Example: multiple RAID employing RAID 0 and RAID 1 to create an array of ten disks. Much as we can define 10 to be 2 * 5 or 5 * 2, we can create our multiple RAID arrays.

RAID 0, then RAID 1

Divide the ten disks into two sets of 5. Turn each set into a RAID 0 array containing five disks, then mirror the two arrays (RAID 1). This is sometimes called a *mirror of stripes*. Called a RAID 01 (RAID zero-one).

RAID 1, then RAID 0

Divide the ten disks into five sets of 2. Turn each set into a RAID 1 array, then stripe across the five mirrored sets. A *stripe of mirrors*. Called a RAID 10 (RAID one-zero or RAID ten).

The standard that *most* of the industry seems to use is that if RAIDlevel X is applied first and then RAID level Y is applied over top of it, that is RAID X+Y or RAID XY or RAID X/Y. Unfortunately, other companies reverse the terms! They might call the RAID 0 and then RAID 1 technique RAID 1/0 or RAID 10.

Some designers use terms RAID 01 and RAID 10 interchangeably.

The result of all this confusion is that you must investigate to determine what exactly a company is implementing when you look at multiple RAID levels.

RAID 0+1: if we stripe together drives 1, 2, 3, 4, and 5 into RAID 0 stripe set A, and drives 6, 7, 8, 9, and 10 into RAID 0 stripe set B. We then mirror A and B using RAID 1. If one drive fails, then the entire stripe set is lost because RAID 0 has no redundancy. The RAID 0+1 array continues because stripe set B still functions. However, reduced to running straight RAID 0 until drive fixed. If another drive in the remaining RAID goes down, you lose data.

RAID 1+0: mirror drives 1 and 2, 3 and 4, 5 and 6, 7 and 8, 9 and 10 to form RAID 1 mirror sets A, B, C, D, and E. Then do a RAID 0 stripe across sets A through E. If one drive fails now, only one mirror set affected. If another drive fails that is in a different mirror pair fails, still fine.

Only two failures in the same mirror set will cause the array to fail, so five drives can fail (as long as they are in different sets) and the array would still be fine.

RAID 1+0 is more robust than RAID 0+1.

If the controller running RAID 0+1 were smart, when a drive failed it would continue striping to the other four drives in its set, and if a drive in a different set failed later it would realize that it could use the analogous drive in the other set since it should have the same data.

This would make RAID 0+1 just as fault-tolerant as RAID 1 + 0.

Most controllers aren't that smart (this is an old slide). In general a controller won't swap drives between component sub-arrays unless the manufacturer of the controller specifically says it will.

The same impact on fault tolerance applies to rebuilding. Consider again the example above. In RAID

0+1, if one drive fails, the data on five hard disks will need to be rebuilt, because the while stripe set will be wiped out.

Disk Array Terminology

Term Definition

Disk or Drive Array Any group of hard drives

Mirroring or shadowing

Striping (RAID 0)

Spanning

RAID

RAID 1

RAID 2 Not used today

RAID 3 Striped data, configured so that each disk I/O operation accesses all drives in parallel, and with non-interleaved parity for redundancy.

RAID 4 Striped data with non-interleaved parity, configured for multi-user systems. Rarely seen today.

RAID 5 Striped data with interleaved parity.

RAID 10 Mirroring (RAID 1) combined with striping (RAID 0)

RAID 10.5 RAID 10, but with fewer drives used for the mirrored backup. This is a cost-effective form of RAID 10. Most commonly, a group of striped drives is mirrored to a single additional drive.

RAID arrays may be controlled either by dedicated hardware (controller-based RAID) or by software (host-based RAID). Hardware RAID is more expensive but provides much better performance and reliability.

RAID Options

Hardware assisted RAID has specific hardware to generate the parity data

The number of inputs and outputs for the video server systems – as well as the data rate for playback and encoding (recording) – are key factors that are restricted by the overall bandwidth of the system. For example, when a system needs playback of 5 streams at 15 Mpbs each, then the bandwidth requirements must be at least 75

Unlike software-based PC codecs (as in Real Video or MS Windows Media Player), video facilities cannot wait for several seconds while a sufficient amount of data is buffered (e.g., stored locally) before playback commences. Any buffering must be continuous and uninterrupted, meaning the data must stream off the drive array in a continuous and uninterrupted fashion.

Parity Calculations (slide)

Outputs XORed together to give us parity. Within a stripe, parity calculated by XOR'ing the stripe units in the stripe and placing the value in the stripe unit storing parity.

RAID 0

Description

Support

Requirements

Centralized Storage

Clip is only stored once and in one location. All users have immediate access to the clip. If the central storage goes down, the whole operation goes down. If a clip is deleted by mistake, it needs to be reloaded from tape. Clips must be moved to the appropriate server.

Distributed Storage

No one server failure can take the whole operation down. A clip may be on several servers, and easily restored if deleted. Generally requires additional storage. The automation system must be able to move clips to the right location and track multiple versions of a clip.

Combination Environment

Good solution. News/production group may have centralized storage for editing/production. When clip done, transferred to playout server.