Experiences with Using a High Quality MPEG Video Server to Aid Classroom Teaching

Mark Gaynor
Harvard University
29 Oxford St.
Cambridge, MA 02138
617 - 496 - 4510
gaynor@eecs.harvard.edu

Vijak Sethaput
Harvard University
29 Oxford St.
Cambridge, MA 02138
617 - 496 - 8020
vijak@eecs.harvard.edu

H.T. Kung
Harvard University
29 Oxford St.
Cambridge, MA 02138
617 - 496 - 6211
htk@eecs.harvard.edu

1. ABSTRACT

In the Spring of 1998 we provided all the lectures of Harvard's network course (CS143) on-line with a high quality MPEG real-time video-on-demand server. Our system allowed students to view lectures using a web based user interface from several different locations on campus 24 hours a day. The students loved the educational value of the video service: it made a complex subject easier to learn; and the availability of the video over the network met the students' busy schedules. This video application provided a real-life example of what we were teaching by demonstrating how network congestion affects the quality of real-time high bandwidth MPEG video sent over a network. Our students found this video service valuable for review, but not a substitute for being present in class.

1.1 Keywords

Network course, video-on-demand server, multi-media education, web based video, distance learning

2. INTRODUCTION

Using a video-on-demand server in the Spring of 1998 as an aid to classroom teaching, we provided our students with access to the complete lectures in Computer Networks taught by Professor H. T. Kung of Harvard University. Our video server system consisted of two server PCs (in the 200+ MHz range), a combined switched/shared Ethernet network, and many client PCs for viewing the content -- it was capable of simultaneously delivering 40 1.5Mbps MPEG (NTSC 352 X 240) streams of video. The video server software is Microsoft's Netshow Theater, previously available free from Microsoft's web site. We found that this solution worked well for our purposes, has a scaleable architecture, and is affordable.

Video-on-demand servers are not new, but high quality MPEG systems like SGI's WebFORCE MediaBase or the Starlight system use custom hardware that is more expensive than we could afford. Our system, built with PCs running Microsoft NT 4.0 server, has storage for 22 90-minute lectures and costs in the $10,000 range. Given our short time frame, limited budget, and current company support contacts, we felt that NetShow was a good option; we were able to set up a prototype system with our current hardware quickly and determine that the product met our demands.

Judging by the feedback from our students, the video-on-demand project was successful. Students liked the web-based access, and started using the service from the first week, before we were completely ready for it. The service continued to become more popular as students began to appreciate the value that it added to the course. We found that the system was particularly useful in helping students review complex topics of the course. We did not find an increase in absences because of the backup the video server provided. The students proved that this video service was more than "neat" technology by continuing to use it all semester.

3. VIDEO-ON-DEMAND SERVERS

A video-on-demand server provides a service that appears to the user as a virtual VCR over a network . It gives you the content you want, when you want it, and how you want to view it. The user must be able to view any title, no matter how popular it is, at any time of day or night. Next, to allow notes to be taken, or information studied, the server should support pausing of the video. Nice features are fast forward and rewind, to help zero-in to the desired place on the video. In addition to the normal VCR controls, a video server should take advantage of the random access abilities of the digital format to scan the video, quickly accessing any section of the content.

Video quality is tied to resolution per video frame, and the number of video frames transmitted per second. Our choice to use high quality MPEG at 30 frames per second precluded the use of Real Audio/Video, a popular video server for Internet and Intranet video applications. We had adequate bandwidth on our network for the 1.5 Mbps required for NTSC MPEG video and felt we needed this quality to allow easy reading of the chalk board.

When compared to video tape, a video server offers several advantages: any client with a network connection possessing enough bandwidth can access the content, and it allows random access to the material. Tape requires the creation and distribution of physical objects, the logistics of which do not scale well. Also, it takes a long time to skip over 90 minutes of video tape. With a video server, once the network infrastructure is in place and the content added, the students can view at will from any client.

Multi-cast of video, like on the Mbone, is a viable way to transmit live events; however, it does not meet the needs of the busy students. Students used the service because of scheduling conflicts, or for review, impossible with multi-casting technology. With the nocturnal habits of many students and the many different ways the students used the system, a multicast service is not a good model for our course intended only for on-campus students.

4. NetShow Theater

The distributed nature of this video server architecture allows building scaleable solutions to reflect the actual demands on the system and meet the budgetary and storage constraints. The minimal system; called an “ice cube”, uses a single PC. This scales easily to a two system configuration, the minimum recommended by Microsoft. From there, additional content servers increase redundancy and capacity. By adding more disks per content server and increasing the number of content servers, the available bandwidth for streaming video increases. The ability to start with a single system as both content and title server and then migrate to a multi-content server configuration allows easy prototyping of a project while maintaining the ability to build up the system as required for more video streams, better redundancy, and increased content.

The title server’s function is to coordinate requests made by the clients into a distributed schedule of videos being served, and to distribute this list to all content servers. To start a video, a client makes a request to the title server. When this new request can be satisfied given the existing schedule of videos being served, the title server then incorporates the new video into the existing schedule. This scheme guarantees that once a stream starts it can continue; however, it does not guarantee how long it will take for the video stream to start, as this waiting time depends on the number of users currently receiving video streams. Once the content servers start the video stream, the title server maintains occasional contact with the client, but is not involved with the transfer of data to the client.

The content servers receive a schedule from the title server, and follow the schedule, guaranteed to be possible given the algorithm [1][2] used by the title server in creating it. The design of the content server enables it to get as much data from the disks as the SCSI bus allows. It then transmits this data on the high speed switched network to the waiting clients, shaping the traffic to provide a steady non-bursty stream of packets on the network. Its configuration should include: at least 96M of RAM, a separate SCSI controller, and set of high speed disks for content storage to achieve the high level of performance required. Connecting each content server with a switched 100Mbps link provides the capacity to serve over 50 MPEG streams per server.

The final component of the Microsoft Theater system is the administrator workstation that manages the entire system, completely independent of its operation. This design allows a lot of flexibility -- any PC can be an administrator workstation, and any administrator server can administer many different video server clusters.

4.1 Why this is a scaleable solution

The design of this video server takes advantage of the inexpensive nature and increasing price performance ratio of PCs today. Many PC configurations have a single SCSI disk attached to a single controller, but the data transfer rate of the disk is typically much lower than the SCSI bus data transfer rate. To maximize the contribution of each content server, the SCSI controller connects to several disk drives allowing a higher utilization of the SCSI bus.[1] This proprietary stripping system also crosses PC boundaries by stripping data across up to 14 separate content servers under control of a single title server. This data stripping allows the title server to devise a conflict-free schedule, assuring that each disk in the system is free to provide the required amount of data within the specified amount of time. By stripping data across several disks on a SCSI controller, and then across several content servers the number of video streams scales to the network capacity.

An attractive feature of this video-on-demand server is the reasonable cost of a system with capacity to hold 22 90-minute lectures (over 25 Gigabytes) and simultaneously serve 40 1.5 MPEG video streams. A system like ours costs about $10,000 (in December of 1997).

[1] But limited by the bandwidth of the 40 MBps ultra wide SCSI controller. Our system used three 7200RPM Segate drives.

5. Our Network

The network used to transport the video traffic is critical to the performance of the entire video server. Real time video of MPEG quality is a demanding service requiring about 1.5Mbps per client with low packet loss and small delays. Given the scaleable nature of this video server, it is not unreasonable to build a system that supports hundreds of streams. These intensive data requirements demand switched high speed networks (either Ethernet or ATM) to have the performance needed to provide good Quality Of Service (QOS) to the clients.

Our network, depicted in Figure 1, shows the local section in Pierce Lab connected with a 100Mbps link to a main distribution switch located in the Science Center. This switch acts as the campus-wide distribution point for the video traffic. Most backbone links in our campus network are switched 100Mbps links that will be upgraded to gigabit links in the near future. We have successfully streamed good quality video from our server in Pierce to a presentation given at the JFK School of Government over a mile away. Our current network configuration provided a good baseline to study what we will need in the future to allow video streaming across campus.

As shown in Figure 1, our network is a combination of switched ports required for the servers along with shared links for the clients used for student viewing. The sharing of client ports on a hub makes good sense given our network setup, but with the price of switched ports dropping we would not build a hubbed system today. We found that sharing the client ports saved us valuable switched ports, and still maintained a good QOS.

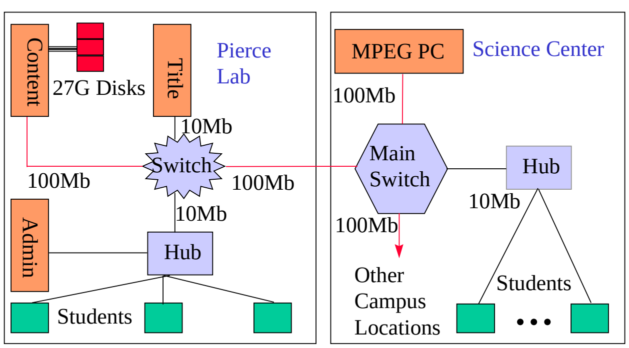

Figure 1 - Our Video Network Capable of supporting 40 1.5 Mbps video streams

6. Content Production

After solving the technical problems, we next worked on improving the video quality by converting the analog video signal from the camcorder directly into the MPEG file in real-time as the lecture was given, as shown in Figure 2 (path b). Originally, we put the content on a VCR tape (Figure 2 path a), and then later re-played the tape into our real time MPEG conversion board[2], not the ideal situation because of the lossy qualities of VCR tape and the extra time required. We later installed a PC with the proper MPEG conversion hardware into the Science Center and fed the live video from the camera directly into the video jacks on the MPEG hardware, thus eliminating the transfer from the VCR, improving the clarity of the video.

There are several options for backup of the content on the video server: VCR video tape, Digital Video (DV) tape, and on-line disk storage. We used both DV tape for its quality, and VCR tape for its ease of use[3]. A better situation would be to set up Netshow to have redundancy in its data stripping, and also keep a copy of each MPEG file on a file server. While on-line disk storage is the most effective, it was too expensive for us, so we used the dual tape backup scheme.

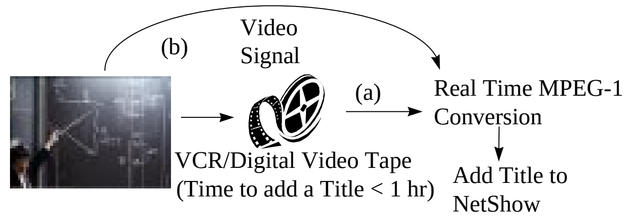

Figure 2 - Our Production Pipeline

7. Student Usage

Our students used the video server for several different reasons, as indicated by a survey and discussion with students : reviewing a complex topic, doing problem sets, studying for tests and watching a missed lecture. Students started to use the server within the first few weeks, and as word of its value spread, increased use throughout the semester. Some students became very consistent, always reviewing the previous lecture right before the next class started. As expected, the amount of use varied greatly from student to student, some relying on the video server for almost every lecture, while others used it only occasionally, if at all.

Whether or not we like to admit it, students will miss classes for many reasons. One surprise was the number of students missing classes for job interviews. Many students interviewed with a different company every week, causing a major disruption in their course work. These seniors were among the heaviest users. We found that in a large class (over 100) we had students watching missed lectures for every single class. Although students did miss class, we did not notice a higher percentage of classes missed per student. In fact, students found the combination of first seeing the in-class lecture, and then reviewing it with the video server provided the most benefit.

[2] Made by Broadway

[3] We did not have a DV tape player.

We like to assume that being in class will always be the preferred method of learning, but we found that for some students this was not true. In particular, one student noted that he paid better attention to the server, saying that, with the instructor so close, he had to pay attention; the student lost track of the fact that it was only a video. By watching several lectures at one sitting, this student also found better continuity in the lectures, and saved time. With large classes, it may turn out that some students concentrate better with the video server because the professor is so close. It is like sitting in the front row, something that we know improves the performance of most students.

Using a high quality MPEG video server to review particularly complex sections of a lecture, or for general review of the material proved invaluable to the students. The high quality offered by NTSC MPEG video allows for reading complex information from the black board, as seen in Figure 3. Understanding complex topics is even more difficult when the lecturer has an accent, or speech impediment. We found that during the more difficult sections of the course, students would watch lectures in small groups, talking about the topics. After our midterm (avg. score 60%) we found that many students used the video server as an improvement tool, re-watching every lecture at least once. It is the random access abilities of the digital video media that are particularly helpful for this use, allowing students to quickly find desired section, or skip over an unwanted section.

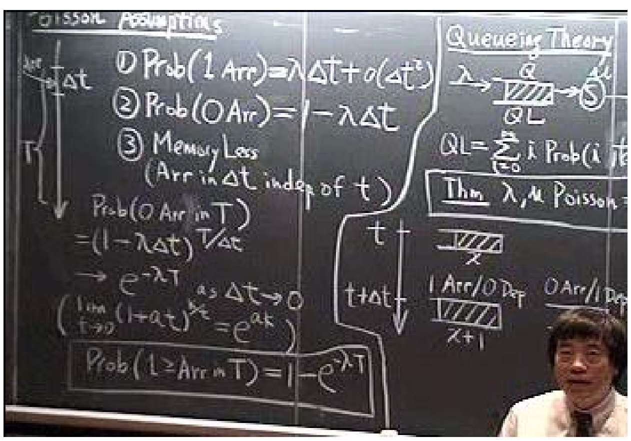

Figure 3 - Content

Having lectures on line allows students to concentrate on watching the lecture the first time, then later, re-watching the lecture while taking careful notes using the pause feature of the video service. We encouraged this by giving out computer generated figures of all board work within a week of the lecture. We found students re-watching the lectures with these notes in hand, adding subtle points brought up in lecture and increasing their understanding of the material.

Unsurprisingly, we found that the heaviest use of the video server was before the midterm and final, or even before a problem set was due. Video-on-demand servers are a great tool for studying, allowing the students to review their points of confusion, very quickly hopping from topic to topic. We found that our client viewing stations were so busy before a test that students formed groups to review the classes. Before a particularly difficult assignment was due, many students would re-watch selected portions of a lecture relating to the problem at hand. This proved to be a great backup to students who thought they had good understanding and adequate notes, but then realized at the last minute while doing the problem set that this was not the case.

A bit of a surprise was that the Teaching Fellows (TFs) also made use of the video service to catch up on a missed lecture or help answer a student's question. It is a great relief for a TF to be able to view a video to verify what happened in class. The TFs found the system easy to use, and most importantly, it gave the TFs an advantage. We could quickly compare what the professor said to what the students heard.

This project exposed our students to a "cutting edge" application -- streaming video at high data rates, showing the theory of what we teach. It is great to talk about what packet loss means, but it is more effective to show students what the effect of losing packets is when watching a video. We invited students to stream video to their dorm rooms, something we knew would only work occasionally. This allowed them to compare viewing from the lab on a fast switched network to viewing from their dorms on a slower shared network. We found that most students enjoyed seeing how network concepts like Quality Of Service (QOS) work in real life.

Overall, most students liked the video service with its easy-to-use web based access, were happy with the quality, and valued the random access abilities of the digital video. The usage patterns showed this as students continued to use the system all semester, something that would not happen if it was not an effective tool. We only had a few complaints. From our point of view, the students were able to get the educational content, and we feel the quality was reasonable in that the glitches, when they occurred, were short and infrequent.

8. Problems

While very happy with our results, we did have our share of problems. These included software glitches, not being able to hear questions from the audience, to network capacity problems, (not a surprise given the Beta status of the Netshow Theater software), and insufficient planning. As with any Beta software, sometimes the cause of the problem is not obvious. The problem reporting feature of our Beta software was particularly poor. Occasionally the system would fail to add a new title, and only provided a cryptic diagnostic message which gave no information as to the cause of the problem. The common fix seemed to be reinstalling all the software, including the NT operating system, not the most practical way to solve a problem. Overall, the number of problems were small and the software worked well for us.

Ultimately, the production of the MPEG is very important and time consuming. We had a dedicated camera person, plus a graduate student who maintained the video server. If lighting, sound levels, camera work, and creation of the MPEG file are low quality then so is the video. It was frustrating that once we ironed out the technical glitches, when we looked at the content, it was sometimes hard to see the instructor's face because of bad lighting. Next time we will be very critical of the factors that account for quality video, since we now know the server can only stream video as good as the original video content.

Real time transport of video traffic is new at Harvard, and part of the motivation for this project was to assess the impact of it on the network. We knew that the network was not up to the task we were asking it to accomplish. We asked students to try streaming video to their dorm rooms, to see what the network would do when overloaded. The parts of the network that had switched bandwidth worked well; it was only when we ventured to the slower shared parts of the network that we started having QOS problems. It did surprise us when the occasional student would successfully watch a lecture from his room.

9. Improvements for next time

We will do this again since we feel it had a positive impact on the course, but realize that improvements to the system will increase the benefit to our students. Some improvements, such as production techniques, are under our control, while others (i.e., network upgrades and sophisticated video indexing) are beyond our control.

Students want access to video services from their dorm rooms, but two problems made this difficult: the network is not up to the task, and many students do not have systems with the CPU power to play the MPEG videos The short-term plan is to distribute video servers and PC clients to each student house, allowing real time video viewing from all on-campus student dorms. Our experience shows that video streaming can enhance the educational experience, and Harvard is building the network infrastructure to allow video streaming campus-wide.

Indexing the video content similar to a book would greatly enhance it educational value. Video indexing is a current topic of research at Harvard and elsewhere. It will be a fundamental improvement, moving digital video into a new era as a tool for education.

10. Conclusion

The video service worked well, and everybody was happy with the results. The students found the web access easy to use and the MPEG video of high quality. Our fear that students would not attend class was unfounded, as most students used the video server for reviewing lectures rather than for making up missed classes. At the end of the semester both students and staff felt the video server added value to our network course.

11. ACKNOWLEDGMENTS

Thanks to Intel for hardware, Microsoft for Beta Netshow software, and many other support staff at Harvard.

12. REFERENCES

[1] W. Bolosky, J. Barrera III, R. Draves, R. Fitzgerald, G. Gibson, M. Jones, S. Levi, N. Myhrvold, R. Rashid. The Tiger Video Fileserver. In Proc. of the Sixth Int. Workshop on Network and Operating System Support for Digital Audio and Video. IEEE Computer Society, Zushi, Japan, April 1996.

[2] W. Bolosky, R. Fitzgerald, John Douceur. Distributed Schedule Management in the Tiger Video Fileserver. OSDI France 97.