Deploying an Infrastructure for Technologically Enhanced Learning

Sami Rollins and Kevin Almeroth
Department of Computer Science
University of California
Santa Barbara, CA, USA 93106-5110
{srollins,almeroth}@cs.ucsb.edu

Abstract: A number of university campuses have undertaken the development of *digital classrooms* that enable presentation of digital media and digital lecture recording. Deploying the infrastructure for a digital classroom is difficult at best even for a technically savvy person. As people from many disciplines become interested in building similar digital classroom spaces, there is a need to produce a useful set of design and implementation guidelines to reduce the project risk and steepness of the deployment curve. The goal of this paper is to report on the experiences we have had in deploying the UCSB digital classroom. The two main contributions of this paper are: (1) a *phased* deployment model; and (2) a discussion of how the proposed technology enables new educational models and techniques.

Introduction

Advances in technology coupled with increased familiarity with technical tools have paved the way for new paradigms in teaching and learning. Instructors are now using media such as PowerPoint slides and digital video in their lectures. Students can take digital notes on laptop computers or Personal Digital Assistants (PDAs). These types of technologies allow students and instructors to communicate digitally across time and space. However, while these tools are readily available, using them in a coherent manner is still a challenge. A number of university campuses have undertaken the goal of developing *digital classrooms* that enable presentation of information using cutting edge multimedia tools as well the capability to digitally record an account of the classroom activity. The account can be used in realtime to enable distance learning or realtime collaboration, or can be archived and reviewed at a later time.

A number of universities have deployed digital classrooms for both teaching and research on new learning methodologies and tools. One of the earliest experiments with this kind of technology was the AT&T Learning/Teaching Theater at the University of Maryland (Schneiderman et al. 95). More recent examples include 405 Soda at UC Berkeley (Wu, Swan, & Rowe 99) and Georgia Tech's eClass (Abowd 99). While the research that has come out of these projects has focused largely on user-level issues, the piece of the puzzle that has yet to be solved is the question of what functionality these classroom spaces should support, and more importantly, how that can be achieved. Without a useful model to draw from, there is an enormous learning curve involved in determining first, what functionality a classroom should support, and second, what technology exists to realize the design. A huge number of tradeoffs need to be considered. It is difficult at best for a technically savvy person to undertake the challenge of deploying a classroom. As people from across disciplines become interested in building similar digital classroom spaces, there is a need to produce a useful set of design and implementation guidelines for ease of deployment.

The goal of this paper is to report on the experiences we have had deploying a digital classroom. By drawing from our experiences, future classroom architects can reduce project risk as well as the steepness of the deployment curve. We identify four classroom functions, and suggest that a classroom should be deployed in four phases corresponding to those functions. In June of 2000 we took on the challenge of deploying a digital classroom at UC Santa Barbara. To date, we have spent approximately $70,000 on our classroom setup broken down as roughly $14,000 for phase 1, $42,000 for phase 2, and $12,000 for phase 3. To date, our phase 4 deployment has leveraged technology purchased in the prior phases. It is difficult if not impossible to deploy a fully functional digital classroom infrastructure before testing or using it. Therefore, it is imperative to support incremental development, deployment, use, evaluation, and modification.

Presentation Facilities

The first phase of classroom infrastructure deployment focuses on providing technology to allow an instructor to give a multimedia presentation in a digital classroom. It is impossible to develop an infrastructure that will accommodate every lecturer or class ever held in a digital classroom. Some professors will use PowerPoint slides while others prefer to use transparencies while still others stick to the standard chalkboard method. In addition, instructors using digital presentation media may require a variety of software. Managing a few pieces of software is tractable, however a system to manage lots of software is not. Fortunately, a large percentage of cases can be accommodated with a standard collection of hardware and software. Minimally, a classroom should include a data projector to show PowerPoint slides or other computer video in addition to providing an Internet connection for a presentation laptop or desktop machine.

Selecting a presentation computer and data projector for purchase requires some thought about the specific classroom and the complete functionality that will eventually be supported by the classroom. We identify three main concerns that need to be addressed when selecting equipment. The first is **compatibility**. A major concern is whether or not each piece of equipment will be compatible with the remaining infrastructure. For example, if the classroom will eventually have a room control system to control various components (e.g., power on/off, input device switching, etc), does the data projector support that type of control? An additional concern with data projectors is how to **install** our mount them. One option is to simply purchase a media cart where a projector can be stored. However, a media cart is not a scalable or permanent solution. The preferred solution is to mount the data projectors into the ceiling. This requires the purchase of a ceiling-mount kit for each projector. Additional concerns include providing a power source as well as ensuring that the ceiling is high enough to mount the projectors out of the way of sight for students and other equipment. Once a data projector and presentation computer have been selected, the next concern is **connecting** them together. The standard solution is to run a VGA cable from the computer to the projector. However, in a digital classroom the distance might be too great thus the quality of the video signal degrades. The solution is to purchase a signal interface -- a device to boost a computer video signal such that it can travel greater distances.

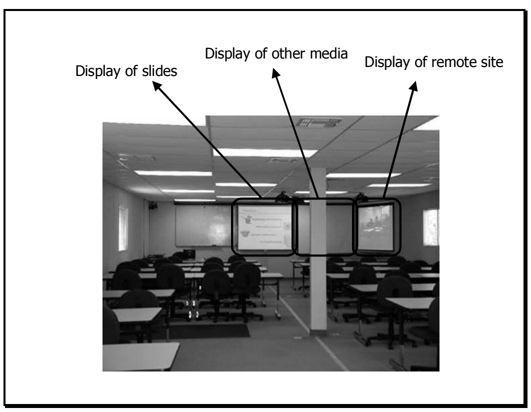

Figure 1: The UCSB Digital Classroom.

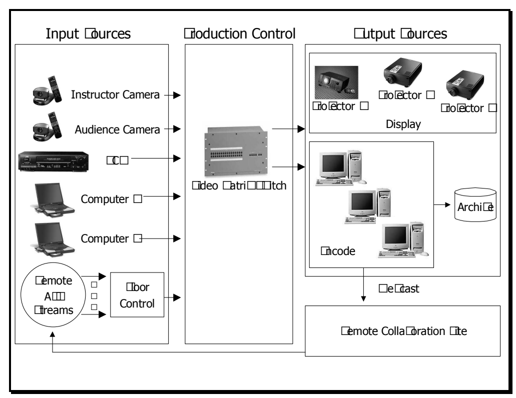

Figure 2: The UCSB digital classroom infrastructure.

The UCSB digital classroom shown in Figure 1 has three ceiling mounted data projectors that project on standard projection screens, two presentation laptop computers with signal interfaces at the front of the room, and one presentation desktop with signal interface at the back of the room. In addition, we provide a VCR for showing standard VHS videotapes. Each computer has a DVD player and an Internet connection and can be used for showing DVDs, presenting PowerPoint slides, and/or web browsing. In addition, speakers may bring their own laptop computers with specialized hardware and/or software and use the data projectors we provide. Finally, students may connect to the Internet using their own laptops or PDAs though a 10Mbps wireless network.

We are currently in the process of identifying other extensions to the infrastructure. Short-term extensions include integrating existing technology such as a document camera. Longer-term extensions include upgrading the display technology to be more sophisticated. Instead of three separate screens we could have a single wall-sized display (Fox et al. 00).

A Webcasting Studio

Once the basic infrastructure is in place to support lectures in a classroom, the next step is expanding the infrastructure to include support for webcasting. A webcasting environment captures audio and video feeds generated in the classroom and sends them, potentially accompanied by other materials such as slides, over the Internet to a remote location. There are two pieces involved in making this happen. The first is to provide support for capturing video and audio of the presentation. The remote audience should be able to see and hear the instructor. Second, multimedia material (e.g., slides) presented to the local audience should also be presented to the remote audience.

The minimal requirements for a webcasting environment are one camera, a microphone for the instructor, and an encoding computer. However, producing a reasonable quality webcast requires a great deal more effort. First, a single camera can be limiting if you hope to capture all of the activity that occurs in a classroom. To capture all of the classroom activity including instructor and student movement, multiple cameras must be mounted in various locations in the room. In addition, capturing slides or a web page can be done simply by focusing a camera on the projection screen. However, capturing the computer video feed straight from the computer can produce a higher quality image. Determining which streams are encoded and webcast at any given time adds a non-trivial bit of complexity to the system. Generally, a *producer* must be available to manually select the stream for webcast. A producer is generally a student or staff member who produces the webcast by controlling the encoding tools and selecting the appropriate video and audio streams.

Managing multiple streams simultaneously introduces a host of complexities. First, devices such as cameras have an associated **control** interface and may be controlled (e.g., zoomed) from a remote control or computer interface. However, when numerous, heterogeneous devices are installed in a room, it becomes difficult to control all of the individual elements. Rather than having multiple interfaces such as computers and remote controls, it is desirable to support a single, integrated interface that supports control of many or all of the devices in the room (Yu et al. 01). Also, it is unlikely that any hardware configuration would support encoding of all possible streams simultaneously. The general protocol is to **select** a subset of all available streams for encoding. In order to accomplish this, the infrastructure must include a video matrix switch. A video matrix switch is a device that takes as input a set of video signals and allows routing of the video signals to one or more of the switch output channels. By routing all video through a video switch, the architecture becomes much more modular. Audio signal capture poses many of the same problems encountered by video stream capture. If an infrastructure supports only a single microphone used by the instructor, the signal can be directly connected to a sound card. However, this model begins to break down relatively quickly. Capturing other audio sources such as audience discussion or the audio track from a VHS video is imperative. The ultimate solution is to install a professional quality audio system that is capable of mixing audio signals from the instructor, microphones placed to capture ubiquitous audience discussion, audio streams from remote sites, and other sources such as video. Finally, video format **compatibility** is a concern. Video capture hardware generally expects a composite video signal. Therefore, the computer video signal sent from a presentation computer cannot be directly encoded as part of a video stream. The solution is to use a *scan converter* to convert the high-quality computer video signal to a composite video signal.

Developing an infrastructure to manage and select multiple media streams in various formats is extremely complex and requires much thought. However, once the architecture to capture a selected stream is in place, the next step is to determine how to encode the audio and video streams into a format that can be easily distributed and viewed by remote participants. There are two primary concerns that need to be addressed. The first is **expense**. Hardware-based encoding solutions provide efficient, high-quality encoding. However, while cheaper solutions are on the way (e.g., NCast - www.ncast.com and VBrick - www.vbrick.com), current solutions can be very expensive. Software-based encoding solutions can also be expensive, but a range of lower-cost solutions exist as well. While **encoding** formats such as MPEG-4 may seem to be the highest quality solution, the chosen encoding format should have a widely available, cross-platform decoder/viewer. Ideally, students would have the viewing tools already installed on their desktop thus avoiding having to download or purchase them. The most common tools currently on the market are RealPlayer and Windows Media Player both of which have freely available and easy-to-install viewers and encoders.

The UCSB digital classroom infrastructure supports webcast of a single video stream selected from a set cameras or other video input. In addition, we have implemented the audio setup suggested by the Access Grid specification (www.accessgrid.org) which includes a number of microphones to pick up ambient sound and a high quality sound mixer[1]. The heart of the UCSB classroom infrastructure is a 12 input, 8 output video matrix switch. The switch supports both composite and computer video. The first two switch inputs (see Figure 2) are the composite video signals generated from the classroom cameras mounted to capture both instructor and audience views. In addition, composite video from the VCR is also routed through the switch. All computer video sources from the presentation machines, as well as remote sources, are also inputs to the switch. However, in order to capture and encode any of these feeds, the feed must be routed through a scan converter and converted from computer video to composite video. The resulting stream is then fed back into the switch and may be selected for encoding. The streams selected for encoding are fed into a video capture card installed on a standard PC. In addition, the mixed audio stream is fed into the sound card of the PC. The audio and video streams are then encoded using either RealMedia format or Windows Media format and webcast to a remote audience.

Remote Collaboration

The one-way distance learning scenario supported by a webcasting studio does not capture the true learning experience. For distance learning to truly be effective, we have to enable remote students to ask questions, participate in discussion, and otherwise appear to be at the local site. This requires two pieces. First, the remote site must have facilities similar to that of the local site. Second, the technology to communicate in realtime between the two sites must be in place.

Ideally, a remote site would be an exact replica of the primary site. In reality though, the infrastructure of a remote site is generally a subset of the infrastructure deployed at a primary site. Minimally, a remote site must include a camera, a microphone, an encoding machine, and a decoding machine. This could take the form of a standard webcam and microphone connected to a student's home PC where the PC is both the encoding and decoding machine. However, if a remote site is designed to support multiple students (e.g., an extension campus) a more complex infrastructure is necessary. For example, displaying video of the instructor on a PC screen is probably not sufficient. In addition, the camera should be able to capture an audience larger than a single person. Fortunately, the problems encountered when developing the remote site infrastructure are the same problems encountered when deploying the primary classroom and thus we can apply the same solutions.

While deploying the remote site infrastructure itself is relatively straightforward and follows directly from the experiences we have already described, deploying a communication layer on top of the infrastructure is more difficult. There are two main topics to be addressed. The first is realtime encoding. Our first inclination was to use the same software encoding solutions for realtime communication that we used for one-way webcasting. However, off-the-shelf encoding software such as RealMedia and Windows Media introduce intolerable delays from 7 to 15 seconds one-way due to buffering requirements. As with many of the problems we have encountered, extremely expensive encoding solutions exist. But, deploying these solutions without knowing whether or not they are going to work is a risky venture. The alternate solution is to use encoding software designed for video conferencing such as *vic* or Microsoft NetMeeting. While off-the-shelf video conferencing software is generally easy to use, quality is sacrificed to meet realtime requirements. Also, in addition to watching video generated at a primary site, students at a remote should be able to ask questions (Malpani & Rowe 97) and access shared components such as whiteboards. Additionally, as the number of remote sites grows, it becomes necessary to manage the sites so that only one remote site is asking a question or sourcing video at a given time. Some of these problems may be solved with standard videoconferencing software. However, the requirements are different for any given infrastructure and differences require specialized solutions.

We have deployed a test remote site we call a *kiosk*. The kiosk has one camera, a single microphone (to be passed from participant to participant), a single encoding machine, and two decoding laptop computers connected to data projectors for display (see Figure 3). We use Microsoft NetMeeting to communicate between the sites. The primary classroom sends slides (using the NetMeeting screen sharing facility), a video feed of the speaker, and an audio stream from the speaker microphone to the kiosk. The kiosk sends a single video stream and a single audio stream back to the primary classroom where it is displayed on a side mounted projection screen. Figure 4 shows the flow of streams between sites. In building the kiosk, we have made a number of simplifying assumptions. First, to avoid the problems of floor control and remote stream selection, we assume

[1] The Access Grid is an initiative to enable research labs and universities to conduct large scale, distributed meetings over the Internet through an *always on* infrastructure.

that there are only two sites participating at any given time. Second, we assume that the only audio stream for either site is generated from the single microphone available at that site thus we do not have to mix the audio signal. Extensions to our infrastructure would include enabling multiple sites to participate at any given time.

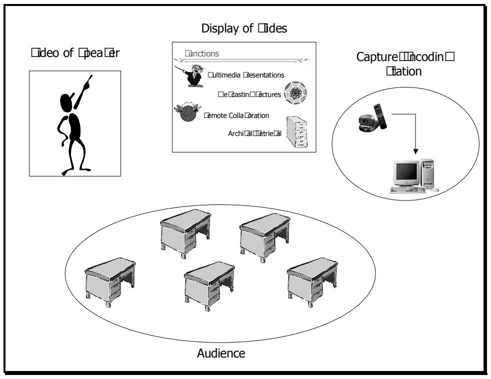

Figure 3: Remote site infrastructure.

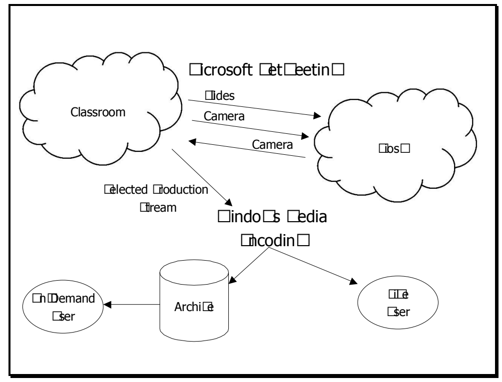

Figure 4: Classroom and remote site communication.

Lecture Replay

One of the primary advantages of recording a lecture or course is that it may be reviewed after the fact. Students may review material at the end of a course or before an examination. In addition, lectures given by guest speakers or experts in their field may be archived and watched by students for years to come (Tsichritzis 99). There are a number of issues involved with recording lectures for replay. However, most of the difficultly lies in providing more functionality over straightforward, sequential playback. The goal of the infrastructure is simply to provide a content base that may be accessed and used to research new methods of access. The challenge is to record and encode the content such that it may be accessed later in different ways using a variety of tools. Support of VCR-style interactivity as well as integration of multiple media types provides a more effective replay experience.

There are two main issues in the deployment of an infrastructure for lecture replay. The first is the deployment of a **media server**. An hour of lecture can be 500Mbytes or more depending on the encoding scheme. Therefore, the first concern is to deploy a media server with enough disk space to hold the recorded lectures. The second issue is to determine which **encoding** standard to use. The simplest solution is to simultaneously save the stream already being encoded for webcast. However, it may be desirable to support postprocessing of the stream such as including synchronization between video and slides (Mukopadhyay 99).

The UCSB digital classroom primarily focuses on replay of the webcast stream. During webcast, the encoded Real or Windows Media stream is saved to a file and may be streamed from the server later. We are still in the early stages of generating a content base and thus have not yet deployed a server to support large amounts of data. We have also built a tool to support synchronization between notes written by participants and the video of the lecture.

Potential Impact on Education

Multimedia in the classroom presents a number of opportunities for students and educators alike. First, using technology, learning can become a more interactive process. Teachers can use a variety of media to teach students in new and different ways while students can use technology such as laptop computers or PDAs to share information and communicate with one another. Moreover, using webcasting and remote collaboration facilities, we can remove the physical barrier imposed by a classroom environment. Enrolled students can “web commute” rather than miss lectures, and students who may have otherwise been unable to take a class at all have more flexibility to choose to attend lecture from their home. Additionally, the number of people who can fit in a room or be physically present no longer limits audience size. For example, if an expert speaker visits a university, the number of people interested in attending her lecture may exceed the capacity of the largest lecture hall on campus. Finally, as webcasting enables a limitless number of people to watch a lecture given by an expert, lecture archival and replay can make the same lecture persistent. Students for years to come can watch and learn from experts in their field

Conclusion

Deploying the infrastructure for a digital classroom is a long and often tedious process. In theory, it involves technical staff, facilities staff, as well as researchers. In the past year, we have brought the UCSB digital classroom online to support the four functions defined by our model to varying degrees of completion. While the process has been slower than we originally anticipated, delays in deployment can be attributed to lack of staff as well as to factors such as back-ordered equipment and equipment incompatibility. However, we are pleased with the resulting infrastructure and its use so far.

To date, we have used the classroom for three different types of events: (1) we have hosted and recorded a standard graduate course; (2) we have conducted a graduate student seminar with participants distributed between the classroom and our remote kiosk; and (3) we webcasted a talk given by a Nobel laureate to an elementary school classroom located in a nearby town. The Nobel laureate's talk was also viewed from other locations in the US, as well as in the UK. Overall, our experiences have been successful. While we are working toward automating the process of set up and configuration as well as lecture production, our infrastructure supports nearly all of the intended functionality.

Throughout the process of designing and implementing our classroom model, we have identified a number of considerations that may not be immediately obvious to the designer or implementer. These considerations range from high-level decisions such as supported functionality to low level choices such as required equipment. We would like to acknowledge many helpful discussions with Larry Rowe and other classroom architects that have helped us to determine the common properties of most digital classrooms. We believe that most digital classroom implementations support similar functionality. Thus, it is unnecessary for each design team to start from ground zero. As campuses around the world begin to embrace technology in their curriculums, it is essential to be able to quickly and easily deploy technologically enhanced meeting spaces.

References

Shneiderman, B., Alavi, M., Norman, K., & Borkowski, E. (1995). Windows of Opportunity in the Electronic Classroom. *Communications of the ACM*, 38 (11), 19-24.

Wu, D., Swan, A., & Rowe, L. (1999) An Internet MBone Broadcast Management System. *Multimedia Computing and Networking*, 1999, San Jose, CA, USA.

Abowd, G. (1999). Classroom 2000: An Experiment with the Instrumentation of a Living Educational Environment. *IBM Systems Journal*, 38 (4), 508-530.

Fox, A., Johanson, B., Hanrahan, P., & Winograd, T. (2000). Integrating Information Appliances into an Interactive Workspace. *IEEE Computer Graphics and Applications*, 20 (3), 54-65.

Yu, T., Wu, D., Mayer-Patel, K., & Rowe, L. (2001). dc: A Live Webcast Control System. Multimedia Computing and Networking, 2001, San Jose, CA, USA.

Malpani, R., & Rowe, L. (1997). Floor Control for Large-Scale MBone Seminars. *ACM Multimedia*, 1997, Seattle, WA, USA.

Tsichritzis, D. (1999). Reengineering the University. *Communications of the ACM*, 42 (6), 93-100.

Mukopadhyay, S., & Smith, B. (1999). Passive Capture and Structuring of Lectures. *ACM Multimedia*, 1999, Orlando, FL, USA.