2007

Virtual Classroom Extension for Effective Distance Education

Radu Dondera

Chun Jia

Voicu Popescu
Purdue University, popescu@cs.purdue.edu

Cristina Nita-Rotaru
Purdue University, crisn@cs.purdue.edu

Melissa Dark
Purdue University, dark@cs.purdue.edu

Report Number:

07-001

VIRTUAL CLASSROOM EXTENSION FOR EFFECTIVE DISTANCE EDUCATION

Radu Dondera
Chun Jia
Voicu Popescu
Cristina Nita-Rotaru
Melissa Dark
Cynthia York

Virtual Classroom Extension for Effective Distance Education

Radu Dondera, Chun Jia, Voicu Popescu,

Cristina Nita-Rotaru, Melissa Dark, and Cynthia York

Purdue University

Abstract

We present the design, implementation, and initial results of a system for remote lecture attendance based on extending on-campus classrooms to accommodate remotely located students. A remote student is modeled with a real-time video sprite. The sprites are integrated into a geometric model that provides a virtual extension of the classroom. The virtual extension is rendered and projected onto the back wall of the classroom. The remote students are displayed at a natural location within the field of view of the instructor, who can conveniently get a sense of their body language and of their facial expression. The system has been deployed in a first classroom and a pilot study indicates that the system promises to deliver quality education remotely. The system relies exclusively on commodity components, therefore it can be deployed in any classroom to allow any course to offer distance education seats.
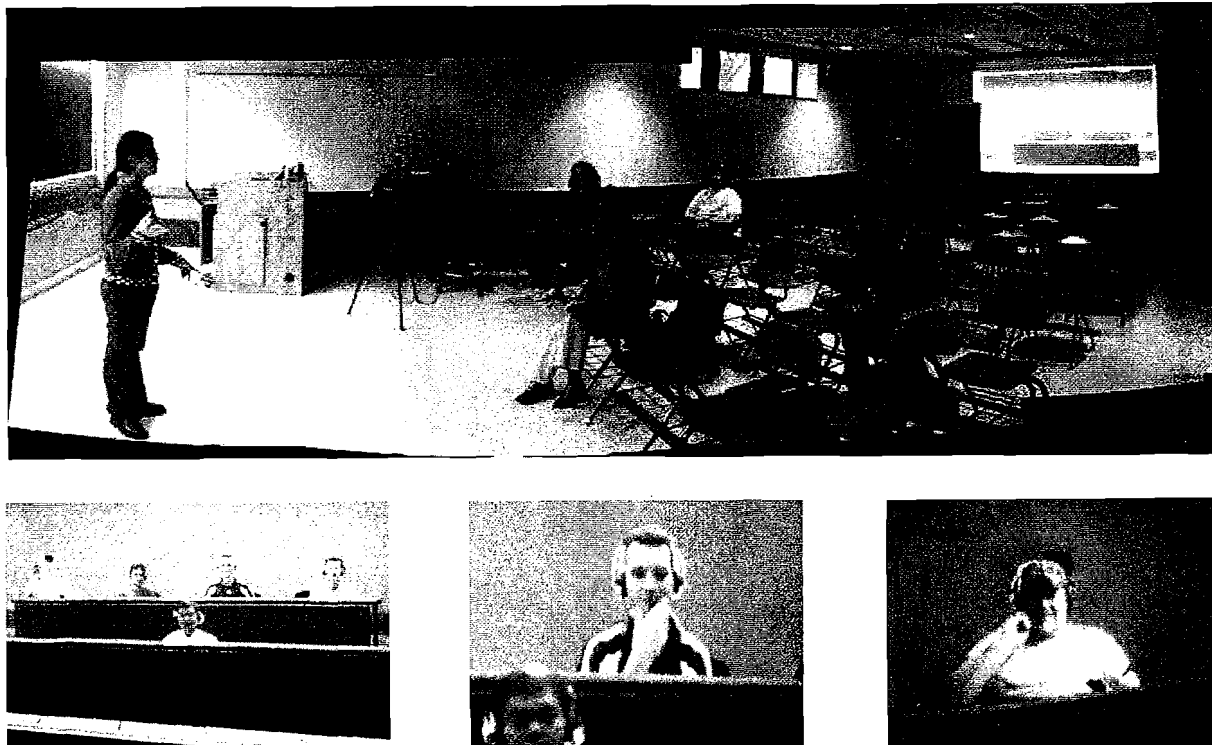
Figure 1. Distance education system deployed in first classroom (*top*), and photographs of the back-wall screen showing remote students integrated into a virtual extension of the classroom (*bottom*).
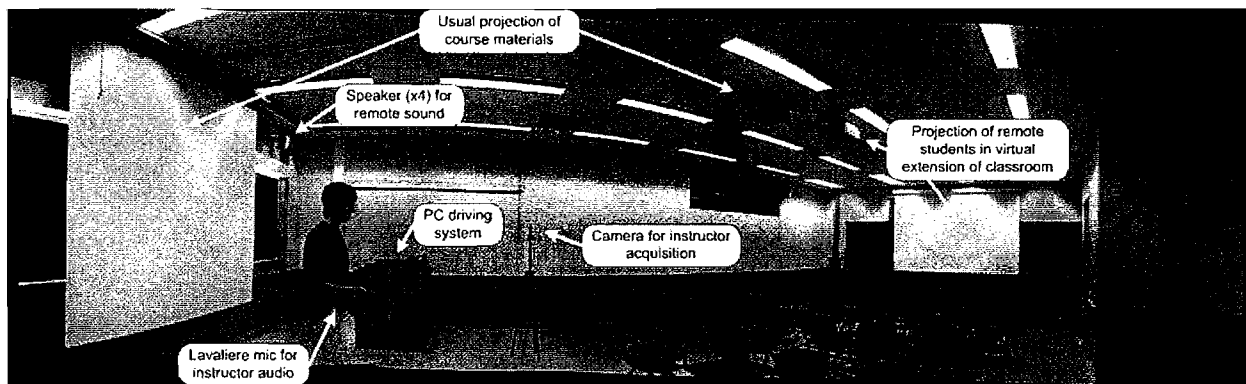
Figure 2. Hardware components for classroom system.

1. Introduction

Distance education services provide access to education for students living in remote areas as well as for working adults, and allow field-authority experts to reach an increasing number of learners. The latest report from the National Center for Educational Statistics indicates that more than half of the higher education institutions in the United States offer distance education services, with enrollment doubling every three years [NCES 2002]. However, current distance education systems fail to match the effectiveness of conventional on-campus education, because of several shortcomings.

First, students engaged in distance education (or for short *remote students*) feel isolated due to lack of interactive communication with the instructor and the on-campus students. Most systems only support asynchronous remote delivery of lectures, and existing synchronous systems provide a low level of interactivity.

Second, remote students do not have access to other proven on-campus education activities such as office hours and study groups.

Finally, the availability of distance education services is limited by the reliance on expensive specialized infrastructure (e.g. teleconferencing-enabled classrooms with attached broadcasting rooms), which requires a large technical support staff. Moreover, the technology places a substantial burden on instructors, who are asked to deviate considerably from their usual approach to educational activity preparation

and delivery. Higher-education institution surveys indicate that cost (program development, equipment acquisition and maintenance), faculty workload, and course quality are factors that prevent to a "major extent" starting or expanding distance education course offerings [NCES 2002].

Research advances in areas such as computer networking, information security, computer graphics, and education science, as well as an almost vertical progress in the performance of commodity audio, video, general computing, networking, and graphics hardware provide new opportunities to reassess the role played by technology in education, and even to develop a new approach to education in general. A debate on whether the current lecture-attendance-based approach to education should be abandoned in favor of a radically new approach, nonetheless interesting, is beyond the goal of this work. Our approach is to leverage technological and research advances in order to increase the effectiveness of distance education to a level comparable to that of conventional on-campus education.

1.1. Contribution

We are developing a distance education system that has unique characteristics that address limitations of current distance education systems. Specifically, our system:

1. *relies on an education-science-inclusive design process* by integrating at all stages expertise and feedback from education science experts, members of our team;

2

2. *provides transparent extension of on-campus educational activities* such as lectures, office hours, and study groups to support distance education, which ensures that remote students are included without significantly altering the preparation or the normal course of the educational activity;

3. *provides support for effective instructor-student and student-student interactions* by using a group communication system as communication platform and by integrating the remote students into a unified virtual environment; this enables the instructor to maintain high-quality visual contact with all remote students simultaneously and thereby to be aware of their facial expression, body language, or desire to interject, in real-time;

4. *can be deployed in any classroom or office* by relying exclusively on inexpensive off-the-shelf components, making distance education an integral part of conventional on-campus education rather than a parallel activity that drains considerable resources;

5. *is evaluated by assessing performance from both the technical and educational standpoints.*

In this paper we report on the design, implementation, and first results of a major component of our distance education system: the remote lecture attendance system. The system has recently been deployed in the first classroom, see top of Figure 1. The operation of the system is illustrated in a video which we ask the reviewers to download from www.cs.purdue.edu/cgvlab/dl/.

Remote students are integrated into a virtual extension of the classroom, which is projected onto the back wall of the classroom. A remote student is acquired with a webcam and is modeled as a real-time video sprite. The sprite is inserted into the geometric model of the classroom extension. Even though each remote student can potentially be located at a different site, the remote students are integrated into a unified virtual environment, which is displayed at a natural location within the field of view of the instructor. The instructor gets a sense of the body language and facial expression of remote students and sees if a remote student raises his/her hand in real time (see bottom row of Figure 1). Audio is captured continually for each remote student and played back in the classroom. Instructor audio and video is continually provided to each remote student.

Our system strives to make distance education an integral but unobtrusive part of conventional on-campus education. This is achieved by relying exclusively on commodity components; the only notable piece of additional equipment is the rear-facing projector (Figure 2). Our system can be deployed in any classroom with minimal cost to allow any course to offer distance education seats. Our system does not imply substantial modifications to the way the instructor prepares and delivers the lecture. The instructor requires little or no training in order to be able to use the system effectively. We foresee that our line of work will ultimately lead to campuses abandoning the current approach of maintaining one or a few distance education classrooms with specialized staff and specially trained instructors, and moving towards merging support for distance education with general audio, video, and IT support.

We conducted a preliminary evaluation of the system through a pilot study involving 5 remote and 15 local students. The goal of the pilot was to measure the effectiveness of the design, independently of networking bottlenecks. Pre- and post-test scores reveal no significant difference between remote and on-campus students. The students also completed a comprehensive survey and participated in focus groups. These indicate that the remote students perceived the ability to naturally interact with the instructor as an important strength of the system, that the remote students would have liked to be able to better interact with the on-campus students, and that the on-campus students were little affected by the distance learning aspect of the class.

We discuss prior work next. Section 3 gives an overview of the system, and sections 4 and 5 describe the classroom and remote student subsystems. Section 6 describes the implementation. Section 7 discusses system assessment. Section 8 concludes and sketches directions for future work.

2. Background

Several systems have been developed that employ audio, video, and communication technology to deliver education remotely (Section 2.1). However, distance education is far from being a solved

problem, as attested by education science research (Section 2.2).

2.1. Prior Distance Education Systems

Several distance education systems have been implemented to date. A virtual classroom multimedia distance learning system [Deshpande and Hwang 2001] extends streaming and conferencing tools by supporting both unicast and multicast networks, by employing a specialized compression algorithm for handwritten text, by improved synchronization and resource allocation for media streaming, and by providing the ability to record the live classroom sessions. Interactivity is limited by rigid communication rules. A single audio and video feed is transmitted at any time, originating either from the instructor or from a remote student. The instructor must grant permission for students to ask questions, which places a strong limitation on interactivity and keeps the focus on the use of technology rather than on the actual educational activity.

A step toward increased interactivity and realism is taken by the Smart Classroom system [Shi et al 2002] which provides a "face-to-face" interactive classroom environment by projecting the images of the remote students onto the wall of a real classroom and by allowing the instructor to move freely in the room. Special network technologies for large-scale access and for heterogeneous network architectures are proposed for real time transmission of data between the instructor and the remote students, but no results are given regarding the effectiveness of these technologies. The system achieves limited interactivity since, as with the previous system, only one student can send his/her audio and video at a time and the instructor has to grant permission. The system projects only the image of the remote student that has the token to speak which does not confer the instructor the ability to maintain contact with all remote students simultaneously, in real time.

A system in which audio and video feeds are sent simultaneously from multiple sites is IRI-h [Maly et al. 2001]. The distinguishing feature of the system is the "shared view", a window with reconfigurable content that is viewed by all participants. The window displays simultaneously a dynamic subset of the participants. However, each participant is shown in a separate sub-window within the shared window, and the lack of a unified visualization reduces the

effectiveness of the visual communication. Moreover, participants are displayed at variable locations within the virtual environment according to the current subset visualized in the shared view, which precludes taking advantage of seating consistency that is known to be beneficial in conventional on-campus learning.

The Virtual Blackboard [Sagias 2002] is a classroom learning environment that allows the instructor and the remote students to manipulate video, audio, text, whiteboard images and 3D graphics simultaneously. The system aims to provide a wide range of modalities for interaction by granting both the remote students and the instructor the opportunity to participate in an online lesson and at the same time exchange information and share applications. Nevertheless, the degree of interaction is limited in that the remote students and the instructor can only communicate via text, whiteboard, email and shared applications, yielding an unnatural interaction.

2.2. Distance vs. On-Campus Education

As distance education has been growing in terms of institutional participation and student enrollment, education science research has been targeted towards assessing current systems and towards drafting guiding principles for putting the effectiveness of distance education at par with that of conventional on-campus education.

Frequently, in distance education remote students do not have the opportunity to interact with other participants directly, or interaction occurs through delayed, asynchronous communication, which precludes in-depth, continuous discussions. Such experiences have been shown to lead to a feeling of alienation and isolation due to the lack of face-to face interaction and support from peers [Galusha 1997]. Numerous research studies that experiment with different strategies for augmenting communication, collaboration, participation, and peer feedback in distance education systems have shown that the level of interactivity is a major factor for increasing learning in distance education environments [Lavooy and Newlin 2003].

Distance education systems are deficient regarding both student-to-student and student-to-instructor interaction [Foshay and Bergeron 2000]. Student-to-instructor interaction is regarded as essential by many educators and students. Research studies on
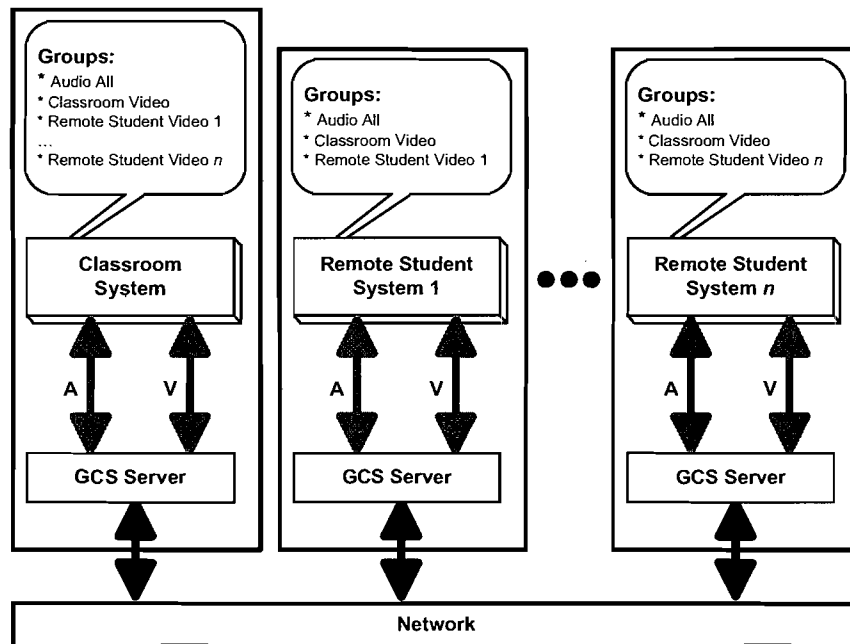
Figure 3. System architecture.

interactivity reveal that students have a real need to make genuine connections with their instructor. Interaction between students and instructors helps ensure the students and instructor develop a feeling of community and connectedness to the course [Lavooy and Newlin 2003].

Distance education systems currently fail to fully engage students because they do not provide a sense of cognitive, social, and physical presence. Cognitive presence refers to having one's cognitive processes focused toward a context or task(s) with one's attention on related stimuli. Social presence, in one aspect, "implies that being with someone virtually feels like being with them physically" [Heeter 2003]. "Physical presence implies being present in (or present to) the virtual or real environment: being there" [Heeter 2003]. A major dimension of presence is interactivity; for example, increased levels of interaction provide more stimuli, which in turn promote a sense of presence.

We conclude that distance education systems can be more effective (increase learning and more likely to motivate students) if they make the remote students

and the instructor feel that the remote students are actually present in the classroom, which should be achieved through high-level of interactivity and visual realism of the virtual environment hosting the remote students.

3. System Architecture

In this section we provide an overview of our system for remote lecture attendance. The system has two main components: a Classroom System that is deployed at the on-campus classroom and a Remote Student System that is deployed for each remotely located student. Figure 3 presents a deployment with n remote students.

The Classroom System extends the classroom to accommodate remote students in a way transparent to the instructor and local students. The image of a virtual 3D room is projected onto the back wall of the classroom to provide additional seats for the remote students. The instructor interacts with the remote students naturally, much the same way she/he interacts with students physically present in the classroom. The Classroom System communicates

5

with each of the sites where remote students are located to send audio and video from the classroom and to receive audio and video from each remote site.

The Remote Student System simulates the classroom environment for a remote student by displaying video of the classroom, and by rendering audio from the classroom and the other remote students in real time. The Remote Student System captures the student with a video sprite which it sends to the classroom in real time. The Remote Student System also captures the student audio which it sends to each of the other remote students and to the classroom in real-time. The audio communication between remote students and the instructor is always on.

The Classroom System and the Remote Student System use a common communication infrastructure for the audio and video transmission. The communication infrastructure is based on a group communication system (GCS). A GCS is a distributed messaging system that enables efficient communication between a set of processes logically organized in groups and communicating via multicast in an asynchronous environment where failures can occur. Services provided by a GCS include group membership as well as reliable and ordered message delivery (e.g. FIFO, causal, or total ordering). The membership service informs all members of a group about the list of currently connected and alive group members, and notifies group members about every group change. A group can change for several reasons. In an idealized fault-free setting, a change can only be caused by members voluntarily joining or leaving the group. In a more realistic environment faults such as processes becoming disconnected or network partitions can prevent members from communicating. When faults are healed, group members can communicate again. These events can also trigger corresponding changes in group membership.

A typical GCS follows an architecture where the major functionality is provided by a set of GCS servers, while applications interact with a server through a GCS client library. An application can be a member of many groups and can act both as a sender and a receiver. Many GCS implement customized group communication protocols relying on UDP unicast and multicast services.

The GCS abstraction provides a modular design and an organization of the communication based on topics of interest to the participants in the system. Groups are just pre-defined names and not reserved IP addresses as in IP multicast. In our case, a single group, *Audio All*, is used for all audio communication since all participants need to receive all audio. All participants subscribe to this group. A participant sends the audio it captures and receives the audio captured by all other participants. Video communication is handled using $n+1$ groups, where n is the number of remote students. We assign one group, *Classroom Video*, for the video captured from the classroom. All Remote Student Systems and the Classroom System are members of this group. Finally, one group is assigned for the video captured from each remote student site. The members of group *Remote Student Video i* are the Remote Student System i and the Classroom System, which send and receive the remote student i video, respectively.

4. Classroom System

The primary task of the Classroom System is to provide support for hosting remote students within a campus classroom, in a manner that interferes minimally with normal lecture activities. We target small to medium classrooms that seat 40 or fewer on-campus students. Such classrooms allow interactive lectures, which most benefit students. Large classrooms preclude all interaction between instructor and students, so augmenting such a classroom with an interactive distance education system is futile. When the lecture is a mere monologue, asynchronous distance education systems do not place remote students at much of a disadvantage compared to on-campus students who can physically attend lecture.

The Classroom System implements the interface between the instructor and remote students. Providing the remote students with instructor audio and video does not pose unusual challenges. However, simulating the presence of the remote students for the instructor is complicated by the fact that each remote student can potentially be located at a different site. Whereas audio contact is straightforwardly enabled by mixing and rendering real-time audio feeds, enabling visual contact is more challenging. A naïve approach would be to display the n video feeds of the n remote students in n separate windows. Such visualization is ineffectual. The instructor cannot
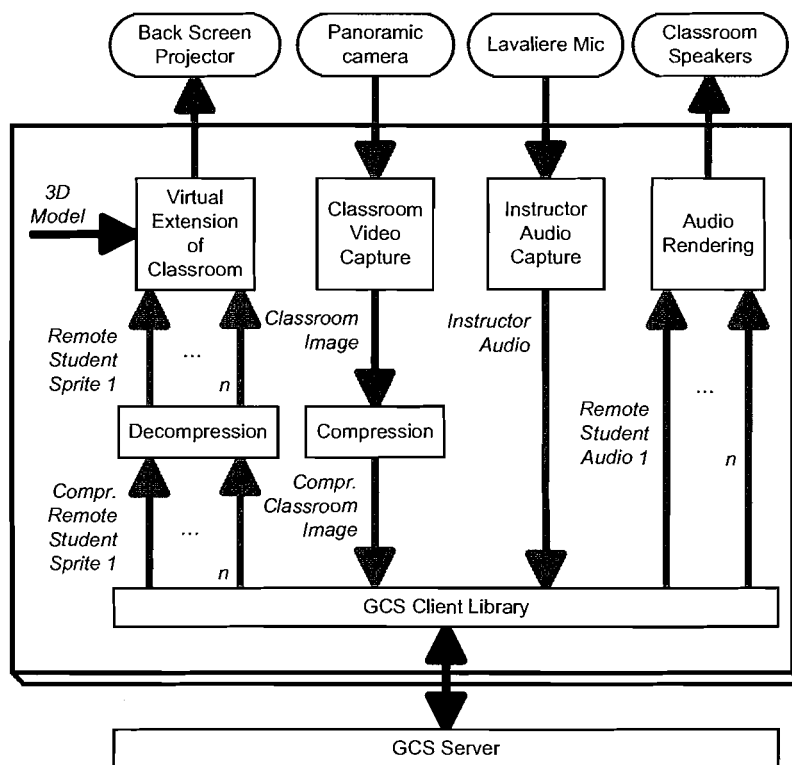
Figure 4. Classroom system.

maintain visual contact simultaneously with all remote students and is forced to scan each window sequentially, spending considerable effort to adapt cognitively to each one of the multitude of contexts.

We have developed a method that allows the instructor to maintain visual contact with the remote students in parallel. The students are extracted from their individual video feeds and inserted into a unified virtual environment, which is displayed at a natural location within the field of view of the instructor (Figure 1). The virtual environment extends the classroom at the back to provide additional seats, which are used by the remote students. The 3D model of the virtual extension is integrated with the images of the remote students and is rendered from a viewpoint that matches the default instructor position. The resulting image is projected onto the back wall of the classroom. The instructor easily monitors the local students and the back wall simultaneously, as if the classroom were larger. The

body language, facial expression, and events such as a remote student raising her/his arm are immediately apparent to the instructor, which enables effective interaction. The penalty of hosting the remote students is essentially limited to the inherent disadvantage of a slightly larger class.

Such an approach has only recently become practical, enabled by advances in video technology (e.g. high-resolution high-frame rate webcams with standard high-bandwidth interfaces, high-resolution high-brightness projectors), in general purpose computing (e.g. PCs that can execute non-trivial per-pixel operations on high-resolution images at interactive rates), and in graphics computing (e.g. add-in graphics cards that can render 3D scenes described with millions of triangles at interactive rates). The classroom system consists of four main modules (Figure 4): Virtual Extension of Classroom, Classroom Video Capture, Instructor Audio Capture, and Audio Rendering, each described in detail below.
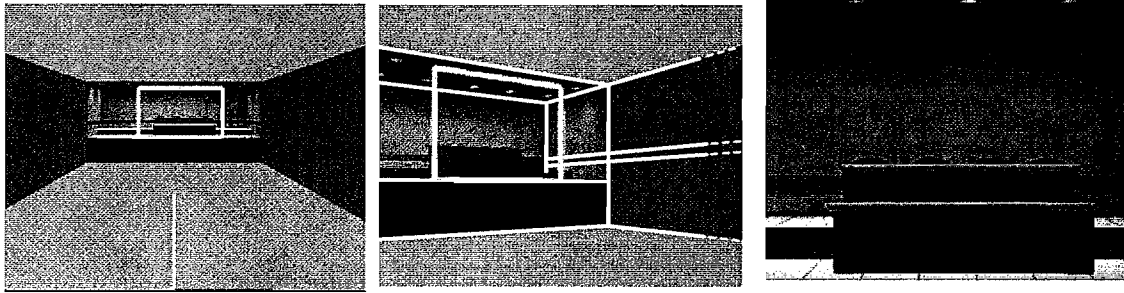
Figure 5. Visualizations of classroom extension (*left* and *middle*) and back-wall projector image.

4.1. Virtual Extension of Classroom

In order to provide a believable virtual environment for hosting remote students, we created a realistic 3D digital model of a classroom that matches the size of the real-world classroom. The side walls, the ceiling, and the imaginary front wall of the virtual extension match the side walls, ceiling, and back wall of the classroom, respectively, as shown in Figure 5. The seating capacity and configuration of the virtual extension is tailored to the number of remote students that it hosts. This not only saves the instructor the disappointment of numerous empty seats, but also has the tangible benefit of a good use of the projected image by spacing the remote students apart horizontally and vertically to prevent that they occlude each other. The virtual extension shown in Figure 1 has two rows of 4 and 3 seats respectively, which is sufficient to host the 5 remote students involved in our experiments.

Realistic large-scale 3D modeling is an open research problem at the confluence of computer vision, computer graphics, and optical engineering. Several approaches are possible. Manual modeling using CAD or animation software produces good results for man-made scenes, but is time consuming and requires artistic talent. Automated modeling based on directly measuring color and depth using acquisition devices such as cameras and laser rangefinders excels at capturing real-world scenes realistically, but suffers from high equipment, time, and operator expertise costs. We experimented with several approaches for producing the 3D model for the virtual extension, including scanning with a time-of-flight laser rangefinder. The manual modeling

approach using state-of-the-art animation software yielded the best results in our case, since the modeler was able to take advantage of object repetition specific to classrooms (desks, chairs). Approximately 100,000 triangles are sufficient to express the geometry of the virtual extension of the classroom, a geometry load that does not pose problems even for low-end graphics cards.

The remote student sprites were placed at predetermined locations in the geometric model to match the seats. The size of the sprites was chosen to approximately match the size of a student actually seated at that respective location. The remote students are updated every time a complete remote student image is received. The virtual extension is rendered 30 times per second and projected onto the back wall of the classroom. In order to achieve this performance yet to provide realism through high-quality lighting, a global illumination solution was pre-computed and burned into the texture maps. This precludes changing the lighting in the virtual extension interactively during the lecture. If such a feature is desired, one could pre-compute high-quality lighting at several intensity levels and interpolate at run time.

4.2. Classroom Video Capture

The Classroom Video Capture module acquires the view of the classroom that includes the instructor and sends it to the remote student sites. A camera with a fisheye lens is used because its large field of view allows keeping the instructor in the frame without having to track the instructor or to move the camera. Although cameras that automatically track the
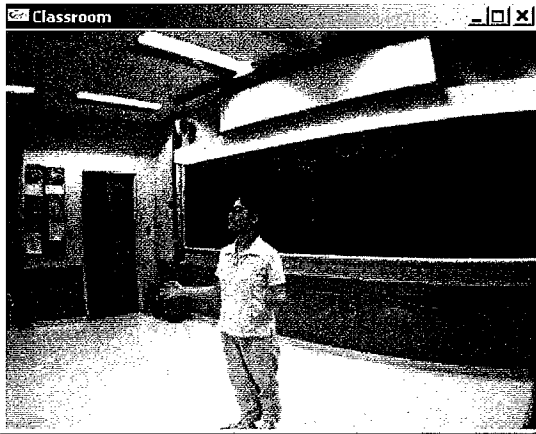
Figure 6. Classroom image delivered to remote students.

instructor are available commercially we have decided against using them because of lack of robustness. Such cameras periodically lose the instructor, which we found to be more disturbing than the lower resolution on the instructor provided by the current solution. The classroom image is down-sampled, compressed, and sent to the remote student sites. Figure 6 shows a typical classroom image as seen by a remote student.

4.3. Audio Capture and Rendering

The task of the Audio Capture module is to acquire audio at the classroom and remote sites. Instructor audio is acquired with a wireless Lavaliere microphone and sent to the remote student sites. The Audio Rendering module receives and mixes audio from each remote student. The mixed sound is rendered on the speakers of the classroom. The same software module is used to render audio for the classroom and the remote students. The software at a location L discards audio packets received from the group that originated from location L.

5. Remote Student System

A Remote Student System i has two tasks. The first task is to simulate the classroom environment for remote student i by receiving and rendering video of the classroom, and by receiving and rendering audio of the classroom and the other remote students. The second task is to capture audio and video of remote student i and to send them to the classroom for integration into the virtual extension. The Remote Student System consists of several modules (Figure 7). The Remote Student Capture module acquires remote student i and builds an effective yet compact visual representation. The Classroom Display module provides remote student i with a view of the classroom through the panoramic images received from the classroom. The Audio Capture module acquires remote student i audio through a headset microphone and sends it to the classroom and the other remote students. Finally, Audio Rendering combines the audio received from the other remote students and from the instructor to provide audio feedback to remote student i via a headset.

5.1. Remote Student Capture

The Remote Student Capture module is a critical component of the Remote Student System providing the instructor with a quality visual depiction of the remote student, in real time. One option is to acquire a 3D model of the student. Compared to modeling the virtual extension of the classroom, modeling the student is complicated by the fact that the scene is dynamic. However, a good approximation of the appearance of the remote student that works well within the virtual extension of the classroom can be obtained if the student is modeled with a real-time video sprite. A video sprite is a video where the background pixels are transparent at all times. The video sprite is a good approximation in our context because the remote student is displayed far away from the instructor (beyond the back wall of the classroom), so the flatness of the sprite is not a concern.

5.2. Background Subtraction

The challenge is to find the foreground pixels in each frame, in real-time. We took the approach of background subtraction. First the remote student moves outside of the field of view of the webcam and a background frame is acquired. Then the background frame is used to determine which pixels changed, which are labeled as foreground. Since the camera adjusts exposure in real time, background pixels will not match exactly the pre-recorded background frame.
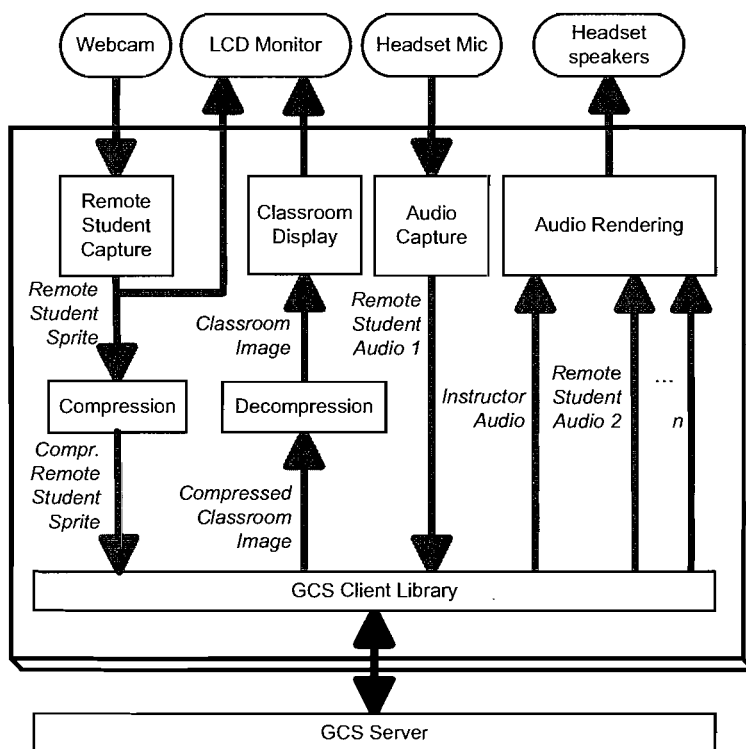
9

Figure 7. Remote student system 1.

Background subtraction has been studied extensively in image and video processing. A good overview of background subtraction techniques was compiled by Piccardi [2004]. Early work in background subtraction algorithms assumed a relatively static scene and used a simple adaptive filter to update the background pixel colors. More recent work utilizes elaborate models of the background to accommodate for changes.

In our context, performance is critical in order to achieve the desired frame rate of 5fps. We developed a simple and efficient background subtraction algorithm that assumes that the background is static. The assumption can be easily satisfied by a student that joins the lecture from a private room. Joining from a large active environment requires the student to choose a location with a static background, as it is the case, for example, when the background is provided by a wall.

Our algorithm performs a series of steps on each new frame: blurring with a flat kernel, correcting the current frame for consistency with the pre-acquired background frame, and computing the foreground pixels as pixels whose color changed significantly. Blurring reduces noise and thus increases robustness. Before comparison with the background frame, the current frame is corrected to take into account the adjustments performed by the camera. Note that although it is possible to disable the dynamic exposure adjustment performed by the camera, such an approach is not desirable since it lowers the quality of the image.

The frame is corrected efficiently using a safety region at the periphery of the frame which is known to be background. In Figure 8, the middle image shows the safety region in magenta. The currently acquired sprite with the magenta safety region is displayed to provide feedback to the student. The student is asked to keep out of the magenta region, and to reacquire the background frame when changes

10

Figure 8. Frame, background, and video sprite compute by background subtraction.

in the background create persistent mismatches visible as misclassified foreground pixels. The sprite is compressed and sent to the classroom.

6. Implementation

We implemented and deployed the system to a first classroom using 5 remote students (Figure 1) placed at different physical locations. The system was implemented using the following hardware (Section 6.1) and software (Section 6.2) components.

6.1. Hardware

The hardware equipment used at the classroom (see Figure 2) consists of

- a PC (3.2GHz Intel Xeon, 3GB RAM, 512MB) with a graphics card (PCIe x16 nVidia Quadro FX 4400),
- a panoramic camera to capture the classroom and the instructor (PointGreyResearch Flea camera with a fisheye lens),
- a Lavaliere microphone to capture instructor audio (Azden 31 LT),
- a rear-facing projector for projecting the virtual extension of the classroom (Hitachi CP-X1250 XGA Multimedia Projector), and
- a sound playback system (Audio Receiver Panasonic SA-HE200K, subwoofer Panasonic SB-WA100, and four speakers Panasonic SB-AFC10).

The classroom was already outfitted with the sound system and a PC for driving the usual front-facing projector; we did not use that PC in order not to interfere with the lectures held in the classroom during the development phase of our system.

At each remote student site the hardware used was a desktop computer, a webcam to capture the student (Logitech QuickCam Pro 4000 webcam), and a headset (Logitech Premium USB Headset 350).

6.2. Software

Our system is developed under the Microsoft Windows platform, which has the advantage of directly supporting any camera, headset, projector, graphics card, microphone, and sound console on the market now or in the foreseeable future, a requirement for the large-scale adoption of our system.

The virtual extension of the classroom is modeled using 3dsmax, a leading commercial animation package. The geometry of the 3dsmax model is saved in the object file format, which is imported into the Classroom System. The virtual classroom extension is then rendered with hardware support using OpenGL.

Video processing in the Remote Student System and audio processing in both the Classroom System and the Remote Student System are implemented with DirectShow filters using threads for concurrency. Concurrency provides a natural separation of independent control flows—aiding development—and also achieves higher performance.

The video processing in the Classroom System is a standalone Windows application and is organized in three concurrent threads: two of them are primarily used to handle the communication, one to send and one to receive, the third is used to handle the graphics

11

component. Specifically, the functionality of each thread is: first thread captures and sends the panoramic view of the classroom, second thread receives the sprite messages sent from the remote students, and third thread renders the virtual classroom into which the remote student sprites are integrated. The video filter in the Remote Student System also uses three concurrent threads, organized in similar fashion: first thread acquires and sends the remote student's video to the network, second thread receives the video feed from the classroom, and third thread displays the classroom. For both systems, the priority of the receiving thread is set to be the highest to guarantee a timely update of the visual feedback to the instructor and to each remote student.

The audio mixer is implemented by a filter that runs concurrently with the video processing as a separate software application. Three concurrent threads are deployed to capture and send, receive, and play back the audio feeds between the remote students and the classroom. Audio is mixed and played back using the DirectSound API. This requires that the header information of the audio sample be sent to the receiving end before any sound can be played back. In order to support students joining the session late, the header is sent periodically. The overhead incurred by repeatedly sending the header information amounts to only 0.02%.

We used the Spread group communication system as our communication infrastructure [Amir and Stanton 1998]. We selected Spread because it provides an easy way to deploy the system while meeting performance requirements, and because it is publicly available. Spread is a general-purpose GCS for wide- and local-area networks. It provides reliable and ordered delivery of messages (e.g. FIFO, causal, or total ordering), as well as a membership service. The system consists of a server and a client library linked with the application. A client can obtain access to the group services by connecting to a server.

Any process—client or server—can fail. If a server fails, all clients connected to that server also fail. When a network partition takes place, Spread servers detect it and continue to provide operation within each connected component. The client and server memberships follow the model of light-weight and heavy-weight groups. This architecture amortizes the cost of expensive distributed protocols, since such protocols are executed only by a relatively small

number of servers and not by all clients. This way, a simple join or leave of a client process translates into a single message instead of a costly full-fledged membership change, which is only triggered by network partitions.

In Spread any group member can be both a sender and a receiver. A client can be a member of many groups. Spread supports a large number of small- to medium-size groups. Each packet in Spread carries its own service type, allowing for fine granularity for the communication. For example, packets sent by applications that require agreement between participants can be sent using a causal or total ordering service depending on the nature of the application while single-source broadcasting applications can use the FIFO service that is more efficient and sufficient for such applications. Spread uses customized protocols relying on UDP unicast and multicast primitives.

7. Assessment

This section gives a quantitative and qualitative assessment of our distance learning system.

7.1. Performance

The main computational tasks of the remote student system are the decompression of the 400x300 classroom frame and the acquisition, construction, and compression of the 160x120 remote student sprite. Each of these tasks is performed at a sustained frame rate of 5 fps, leveraging efficient MPEG-4 compression and decompression implementations provided by the XviD library, and the assumption of a static background for sprite construction.

The classroom system compresses the classroom image, decompresses the sprites of the remote students and updates their respective textures in the 3D virtual extension of the classroom. Five remote students are comfortably handled by the single classroom workstation at 5fps. In fact, simulations show that decompression performance and bandwidth to the graphics card do not become a factor up to 50 remote students (modeled with 5fps 160x120pixel sprites).

The average bit rate for sending and receiving a 160x120 remote student video sprite at 5fps is approximately 35 kbps; the average bit rate of the classroom video is approximately 355 kbps for

400x300 images at 5 fps. The audio bandwidth requirement is negligible by comparison. This implies an upload/download bandwidth requirement of 35/355 kbps at each remote student site, which is within the limits of commodity broadband networking (e.g. DSL or cable modem). For the classroom system the required upload/download requirement is 355/175 kbps, which is within the limits of the connectivity available at most educational institutions.

7.2. Informal assessment

During a demonstration of the system, the authors, including education science experts, had the opportunity to informally assess the system. It was found that the system provides a quality audio and video link between the instructor and the remote students.

The quality of the sound is excellent, allowing for vocal intonations and nuances to be detected. This facilitates a sense of rapport among the remote students and the instructor. It also partially supports a sense of rapport between local students and remote students; the remote students are somewhat disadvantaged in that they can hear the instructor, but not the local students.

The instructor interacts naturally with both local and remote students as she/he can hear and see both sets of learners. The remote students can see the instructor clearly. The instructor is able to move freely while lecturing, which is an important advantage over systems for distance lecture delivery that employ a fixed camera and have the instructor place bound.

The sense of the remote students being present seemed greater for the instructor than for the remote students themselves, because while the local participants could see the remote students, the remote students could not see themselves integrated into the classroom. The quality of the remote students' image is sufficiently high for the instructor to discern whether the remote students are confused or disengaged.

The remote student can raise her/his hand to ask or answer a question. The remote student has no control over what they can see, they are restricted to the panoramic image of the classroom received. The instructor and local students have control in that they can choose whom to look at: local students, the instructor, or the remote students. Presence and interaction could be increased by providing the remote students with a view of the entire classroom that includes the remote students. However, this system shows promise for enabling interaction and a sense of presence not possible with current systems.

7.3. Pilot study

In order to further assess the effectiveness of the design of the distance learning system we conducted a pilot study involving 15 local and 5 remote students. None of the students had been involved in a distance learning class prior to the pilot study.

The remote students were distributed throughout the building of the local classroom and the remote student workstations were connected to the local classroom through a 100Mbps local area network. This eliminated confounding factors such as occasional networking bottlenecks, and also enabled our team to observe the remote students during the pilot.

The students attended 4 lecture sessions, each lasting 1 hour. The lecture topics were introduction to digital video, camera operations, video formats, and video delivery. Before the first session the remote students were briefed on how to use the system (e.g. how to turn audio communication on and off, and how to acquire the background to be used for the construction of their sprite).

The students were tested before and after the 4 lectures. The tests show no significant difference in pretest scores between the local and remote students, which indicates that the local and remote groups were roughly equivalent. Although the student group sizes were small, there were no outliers, which strengthens our confidence in this conclusion. The posttest scores were higher for each group than their pretest scores, which indicates that both groups learned. The posttest scores also show no significant difference between the two groups, which suggests that both groups learned the same amount.

After the 4 sessions, in addition to the posttest, the students also completed a comprehensive survey comprising 93 questions and participated in focus group sessions. A detailed description and analysis of the survey and focus groups is beyond the scope of this paper. The conclusions we have drawn are:

- According to the remote students, a strength of the system was the ability to directly talk to the instructor and to other remote students.
- According to the remote students, a weakness of the system was the inability to see, hear, or talk to local students.
- The local students were little affected by the involvement of the remote students, and, in general, by the fact that the lecture was also delivered remotely.

The instructor that delivered the lectures during the pilot has extensive experience with prior synchronous and asynchronous distance learning systems. During an interview after the pilot he indicated that the distance learning system provides a closer connection to the remote students than possible with prior systems. He noted that he was able to talk to the remote students frequently and casually, and that, overall, the lectures "did not feel like distance learning".

8. Conclusions and Future Work

We have designed, implemented, and deployed a distance education system based on integrating the remote students in a unified, highly-interactive virtual environment. Most classrooms already have a front-facing projector driven by a PC—the main piece of additional equipment is a projector facing towards the back of the classroom. The projector paints a virtual hole into the back wall which reveals a virtual extension of the classroom that hosts remotely located students. Preliminary tests indicate that the configuration is effective. The remote students feel and are perceived as being present in the classroom.

We are currently augmenting the system with capabilities such as instructor tracking, morphing of classroom extension to smoothly adjust to the number of students, exaggerated visualization of remote students for improved interaction efficiency, electronic whiteboard support, and distributed class materials. The deployment of the first prototype is an important milestone that enables future work in several major directions.

One direction is the use of the system in the context of actual courses to complete a rigorous evaluation of the educational impact. A second direction is to deploy the system over the Internet, which brings

new challenges such as addressing security and privacy issues.

Most of the current effort concentrated on improving the instructor-side of the interface between instructor and remote students, since it is presently the most deficient thus the most critical. A third direction of future work is to research improving the interface on the remote student side, by possibly letting the student see her/himself integrated in the classroom, together with the other remote students.

Finally, we will develop support for other on-campus interactive educational activities such as attending office hours and study groups. In addition to researching the best way of arranging the remote student avatars such as to enable round-table-like discussions and to add the capability to freely exchange hand-written notes, supporting study groups effectively will require creating a comfortable virtual setting conducive to productive collaboration. One important sub-problem is producing highly-realistic digital 3D models of coffee shops, student union lounges and other preferred campus venues.

9. Acknowledgments

References

Amir, Y. and Stanton, J. 1998. The Spread wide area group communication system. In *Technical Report CNDS-98-4*, Johns Hopkins University, Department of Computer Science.

Deshpande, S. G., Hwang, J.-N. 2001. A real-time interactive virtual classroom multimedia distance learning system. In IEEE *Transaction on Multimedia*, vol. 3, no. 4.

Foshay, R. and Bergeron, C. 2000. Web-based education: A reality check. In *TechTrends*, 44, 16-9.

Galusha, J. M. 1997. Barriers to learning in distance education: University of Southern

Mississippi". Retrieved May 10, 2006 from http://www.infrastruction.com/barriers.htm

Heeter, C. 2003. Reflections on real presence by a virtual person. In *Presence: Teleoperators and Virtual Environments*, 12(4), pp. 335-45.

Lavooy, M. J. and Newlin, M. H. 2003. Computer mediated communication: Online instruction and interactivity. In *Journal of Interactive Learning Research*, 14(2), pp. 157-65.

Maly, R. et al. 2001. IRI-H, a Java-based distance education system: its architecture and performance. In *Journal on Education Resources in Computing (JERIC)*, vol. 1, issue 1es.

NCES (2002) Post-secondary education quick information system. In *Survey on Distance Education at Higher Education Institutions: 2000-2001*.

Piccardi, M. 2004. Background subtraction techniques: A review. In *Systems, Man and Cybernetics, 2004 IEEE International Conference on*, vol. 4, pp. 3099-3104.

Sagias, Yiannis N. 2002. Distributed and online distance lecturing environment (the Virtual Blackboard Project). In *Educational Technology & Society*, 5 (1).

Shi, Y., Xie, W., Xu, G. 2002. Smart remote classroom: Creating a revolutionary real-time interactive distance learning system. In International *Conference on Web-Based Learning*.