Asynchronous Collaboration Around Multimedia and Its Application to On-Demand Training

David Bargeron, Anoop Gupta, Jonathan Grudin, Elizabeth Sanocki, Francis Li[1]
Collaboration and Multimedia Systems Group
Microsoft Research, Redmond, WA 98052
{davemb, anoop, jgrudin, a-elisan}@microsoft.com, fli@cs.berkeley.edu

Abstract

Multimedia content is a central component of on-demand training and education delivered over the World Wide Web. Supporting asynchronous collaboration around educational multimedia is a key requirement for making the online educational environment effective. A multimedia annotation system tightly integrated with email provides a powerful platform on which to base such functionality. Building on our earlier work with multimedia annotations [3], we present user interface and system extensions to support asynchronous collaboration for on-demand training. We report results from a study of the effectiveness of the system, including student experience, instructor experience, and user interface appropriateness. Overall, the student experience was positive: they appreciated the flexibility of on-demand delivery and benefited from the collaborative features.

1. Introduction

Many see the Web as an irresistible platform for education and workplace training [4][16], with multimedia content available on-demand. Educators foresee vast improvements in cooperative inquiry. Students anticipate convenience and access to education that was unimaginable just a few years ago. And universities and corporations envision lower costs and increased efficiency. However, effective on-demand learning faces significant challenges.

Traditional face-to-face classroom instruction enhances the material in books by providing an instructor's perspective and by promoting discussion. In-class discussion takes place in the context of the lecture material. The learning is reasonably effective (we are all products of it!), and forms the base case to which other alternatives must be compared.

One of the early alternatives to this model explored the inclusion of remote students. Lectures were broadcast via TV to remote students' sites and telephone back-channels were provided for questions to the instructor. Stanford University launched such a program over 25 years ago, for example, and it is still very popular and profitable today. Hundreds of studies have shown that students can learn as much from such broadcast lectures as from live classroom attendance [34]. However, a significant drawback is that this model is "same-time:" Everyone must meet at an appointed time and date. Students cannot participate on-demand, and the model is not very scalable. Broadcasting a lecture simultaneously to thousands of students would fundamentally alter the interaction model and thus the learning experience.

Another model, built around the distribution of videotaped lectures and supporting materials to students for on-demand viewing, has been shown to produce learning comparable to live classes [34]. But this relatively low-tech approach suffers from higher-than-normal dropout rates and greater barriers to interaction.

Today, however, advances in technology support on-demand delivery of multimedia over intranets, and, increasingly, over the Internet. This offers new possibilities for enhancing the on-demand educational experience, including new opportunities for asynchronous interaction.

For example, the Internet is allowing students to watch courses on-demand using streaming media from their homes and offices. And students who watch lectures online (e.g., [27]) can often participate in discussions around the material using email, newsgroups, and chat rooms. But even so, the interactions occur outside the lecture context. Class bulletin boards or newsgroups are usually accessed after, not while, a student listens to a lecture. And the commentary that is generated is not tightly linked to the relevant points in the lecture.

As on-demand multimedia educational content becomes more sophisticated, preserving the context of comments and questions, and providing enhanced means of interaction among students and teachers, grows ever more important. This paper describes a system built to address these issues.

In this paper we describe extensions to MRAS (*Microsoft Research Annotation System*) [3], our general

[1] Current address: Computer Science Department, University of California Berkeley, Berkeley, CA 94720

architecture for multimedia annotations, and how we employed them to address short-comings in on-demand educational systems. We built a set of closely integrated, independently reusable client components that serve as a flexible platform for asynchronous collaboration. All of the components were designed to be web-based and programmable so they could be embedded and controlled in web pages. We collected user requirements, then designed and built a new on-demand education interface using the components. Finally, we studied the effectiveness of the new interface during two on-demand offerings of a popular introductory training course.

The remainder of the paper is organized as follows. The scenario in the next section conveys the activity we envision supporting and the functionality we designed to do so. We then briefly discuss related work. Next, we give an overview of multimedia annotations and how MRAS supports them. We then describe the MRAS extensions intended to better support asynchronous collaboration and workplace training. Finally, we describe the study design and findings, including student and instructor experience and the appropriateness of the user interface.

2. Scenario

A student logs-in to watch a lecture at 10pm from her home computer. Through her web browser she receives the audio-video of the professor, the associated slides that flip in synchrony with the video, and the notes associated with the slides. In addition to typical VCR-like navigation features for the lecture video, there is a table of contents of slide titles, and clicking an entry jumps or "seeks" the presentation to the appropriate slide and audio-video.

The key additional innovation over systems available today is that the student also sees questions and comments entered by classmates who watched the lecture before her, as well as responses from other students, the teaching assistant (TA), and the professor. These questions are linked to the lecture content, and as she watches a lecture, questions asked during that portion of the lecture are automatically highlighted or "tracked." The content of a question appears in a preview window, and if one piques her interest she can seek the presentation to it. As she is watching, she sees a question that nobody has answered yet. She chooses to type a response, which is automatically registered with the system and displayed with the question. The person who posed the question is notified of the reply by email.

Later, the student has a question. She selects the "ask question" button, then types a subject header and her question. Afraid that the question may sound uninformed, she makes it anonymous. In addition, she enters the email address of a friend, who may be able to answer it before the TA gets to it. When she saves the question, it is added to a pre-existing shared "discussion" collection and is automatically emailed to the TA alias and to her friend. A TA browsing through his email sees the question arrive and opens the message. The email includes the text of the question along with a URL pointer to the point in the lecture where the question was asked. It also contains enough meta information for a reply to be added to the annotation database, making it visible to students who later watch the lecture.

The student can similarly record personal notes, also linked to the lecture. These are added into a different collection, with permissions set by the student.

This scenario suggests how asynchronous environments can enjoy many of the benefits of the question-and-answer and discussion that occurs in "live" classrooms. Our multimedia annotation system, MRAS, is designed to support this scenario by implementing multimedia annotations, a fine-grained access control structure, and close integration with email.

3. Related Work

Annotations for personal and collaborative use have been studied in several domains. A number of annotation systems have been built and studied in educational contexts, primarily focused on personal annotation. CoNotes [8] and Animal Landlord [25] support guided pedagogical annotation experiences. Neither focused on multimedia lecture scenarios, and their functionality is not as general or rich as MRAS (e.g., tight integration with email). Some studies of handwritten notes [17] have shown that annotations made in books can be valuable to subsequent users, the benefit we hope to extend to video content.

The Classroom 2000 project [1] focuses on capturing all aspects of a live classroom experience (including whiteboard strokes) and making it available for subsequent student access. The same is being done, with less rich indices, by most major universities exploring the distance learning market [27][30][7]. Although MRAS is a powerful system for storing indices, our focus is on more dynamic, asynchronous content.

WebCT [32] and Lotus LearningSpace [15] are commercially available systems for creating educational web sites. They support chat, email, and bulletin-boards for communication, and a degree of association between the artifacts of communication (e.g., email messages) and the context in which they were created (e.g., a particular web page). Neither product offers MRAS's rich support for multimedia, for video annotations, and for fine-grained organization and sharing of annotations.

AnswerGarden [2] and Organik [21] support the collection of questions and answers in indexed, searchable "FAQ" databases. Both are integrated with email to route questions to the most appropriate expert. These and other systems [18] can provide high quality access to information; however, they do not generally maintain a connection between the information in the database and the context in which it was created. Unlike MRAS, they do not consider questions to be annotations which have meaning and significance in a specific context.

The Vicarious Learner Project [31] explored several ideas related to our work and demonstrated that student dialogs can be a valuable learning resource for subsequent students [19]. The Computer Supported Intentional Learning Environments (CSILE) project [6] and its commercial outgrowth called Knowledge Forum have shown that collaborative learning can enhance the depth of each students' understanding of the subject matter [24]. Our present work takes this a step further by tying student dialogs to the context in which they occurred.

The MRAS system architecture is related to several designs. OSF [25] and NCSA [11] have proposed scalable Web-based architectures for sharing annotations on web pages. ThirdVoice [28], NovaWiz [20], Hypernix [9], uTok [29], and Zadu [35] have all recently released (or are about to release) commercial web note-sharing systems. These are all similar in principal to MRAS, however none supports fine-grained access control, annotation grouping, video annotations, or rich annotation positioning. The web-based Knowledge Weasel [12] offers a common annotation record format, annotation grouping, and fine-grained annotation retrieval, but does not support access control, and it stores meta data in a distributed file system, not in a relational database. The ComMentor architecture [23] is similar to MRAS, but access control is weak and annotations of video are not supported. And to the best of our knowledge, no significant deployment studies have been reported for any of these systems.

Considerable work on video annotation has focused on indexing video for video databases. Examples include Lee's hybrid approach [14], Marquee [33], VIRON [10], and VANE [5]. They run the gamut from fully manual to fully automated systems. In contrast to MRAS, they are not designed as collaborative tools for learning and communication.

4. MRAS

This section gives a brief overview of multimedia annotations, the MRAS base infrastructure, and the first-generation MRAS user interface.

4.1 Multimedia Annotations

Multimedia annotations, like notes in the margins of a book, are meta-data associated with multimedia content. There are a few additional unique aspects when we consider them in the context of audio-video content and client-server systems:

Annotations are anchored to a point (or a range of time) in the timeline of a video presentation, rather than to points or regions on a page of text.

Annotations are stored external to the content (e.g., audio-video file) in a separate store. This is critical as it allows third parties to add annotations without having write access to the content. Students should not, for example, be able to modify the original lecture.

Because annotations persist in a database across multiple sessions, they are a fitting platform for asynchronous collaboration, where users are separated in time. Furthermore, with appropriate organizational and access control features, they allow for structured viewing and controlled sharing among users (for example, they can be grouped into personal notes and public question-and-answer collections). Finally, they enhance the end-user experience by being displayed "in-context," i.e., at the anchor point where they were made.

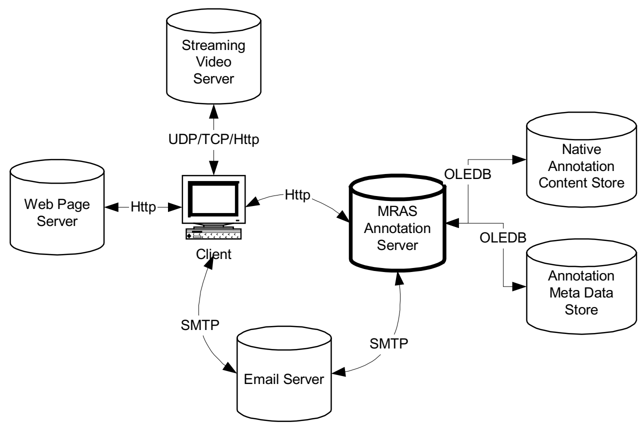

Figure 1: MRAS System Overview. The MRAS Annotation Server fits into a standard multimedia network architecture.

4.2 MRAS System Overview

The MRAS prototype system is designed to support annotation of multimedia content on the web. When a user accesses a web page containing video, the web browser contacts the web server to get the HTML page and the video server to get the video content. Annotations associated with the video on the web page can be retrieved by the client from the Annotation Server.

Figure 1 shows the interaction of these networked components. The Annotation Server manages the Annotation Meta Data Store and the Native Annotation Content Store, and communicates with clients via HTTP. Meta data about multimedia content are keyed on the content's URL. The MRAS Server communicates with Email Servers via SMTP, and can send and receive annotations in email.

4.3 Original User Interface

Part of the original MRAS User Interface (UI) is embedded in the web browser and part is external with separate windows [3]. Correspondingly, Figure 2 shows the MRAS toolbar embedded at the base of the browser window, and the MRAS "View Annotations" window in the foreground. The toolbar is used to specify which annotation server to connect to and what annotation-sets (questions, personal notes, table-of-contents, etc) to retrieve, and for performing "global" operations such as adding a new annotation.

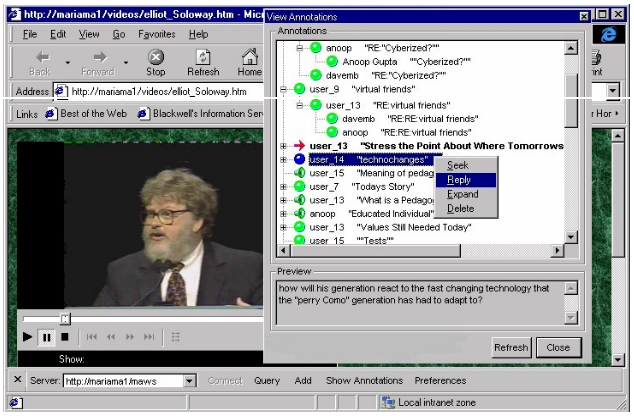

Figure 2: Original MRAS User Interface. This UI is useful for browsing the web, but requires too much input from the user.

Once annotations are retrieved from the server, their headers (e.g., author and subject fields) are displayed in the overlaid "View Annotations" window in video-timeline order. They can be edited or deleted, and replied to (thus forming threaded discussions). They can also be used to navigate within the video presentation. The annotation closest to the current time in the video is highlighted by a red arrow, thus keeping the user's view synchronized with the video. The content corresponding to it is displayed in a preview pane at the bottom of the window.

5. Extending the User Interface

Although the original MRAS UI works well for some tasks, it has several weaknesses when it comes to the lecture-based on-demand education classes we wanted to support. This section describes our analysis of usage requirements for on-demand education incorporating multimedia, and the user interface we developed as a result.

5.1 User and Task Requirements

To design the new interface, we monitored and talked to several instructors and students in corporate training classes and drew upon our own experience as students and instructors. We reached the following set of requirements for support of on-demand training:

Simplify the annotation user interface: The original MRAS interface requires too many decisions from users, including which annotation server to connect to, which annotation sets to retrieve, and which annotation set to add new annotations to. In the new interface, the author of the class web pages makes all such decisions; they are invisible to the students. Less technologically sophisticated students can therefore participate more easily.

Integrate annotation features directly into the on-demand education interface: The original MRAS interface does not allow annotations (headers or content) to be embedded directly into a frame within the web browser. The "View Annotations" window frequently interfered with the content underneath it, for example by covering lecture slides (see Figure 2). Users are sometimes confused by this interference and their attention is often split. Clearly, we had to extend the design so that UI components could be embedded within regular browser frames.

Visually categorize annotations: When annotations from multiple annotation sets (i.e., user-defined, access-controlled collections of annotations) are retrieved in the original MRAS interface, they are all displayed in the same "View Annotations" window. Mixing of annotations is not always desirable, and students wanted to keep their personal notes, shared discussions, and table of contents annotations in separate windows to avoid clutter.

Always track annotations: In the original interface, students can choose to have annotations "track" the video, highlighting the nearest annotation(s) to the current point on the video timeline. This feature was popular enough to become the default.

Automatically email annotations to a class distribution list: The original MRAS UI supports emailing annotations, however unless an annotation includes an email address explicitly, it can only be seen through the MRAS UI. Instructors wanted to be maximally responsive to student questions and comments, so we had all public annotations sent by default to the instructor, and also to other class members on a class email distribution list. As before, personal annotations have an optional email address field, and replies to emailed annotations are added to the set of the annotation they are replying to.

Display instructor notes along with slides: When listening to a lecture and watching corresponding slides (see Figure 3), students wanted access to "instructor's notes" at the same time. We needed to display them along with the slides and have them change as the slides flip in synchrony with the video.

Streamline the use of annotations as a navigational tool: The students wanted to be able to browse the various annotations others had made while watching the lecture. They wanted to use this browsing capability to decide which portions of lecture to watch.

Integrate demos into the multimedia educational content: The instructors wanted to show short, high-resolution video demos of examples given in the live classroom. These are not part of the main lecture video, so when it is time to show a demo, the main lecture video should automatically stop, the demo video should start

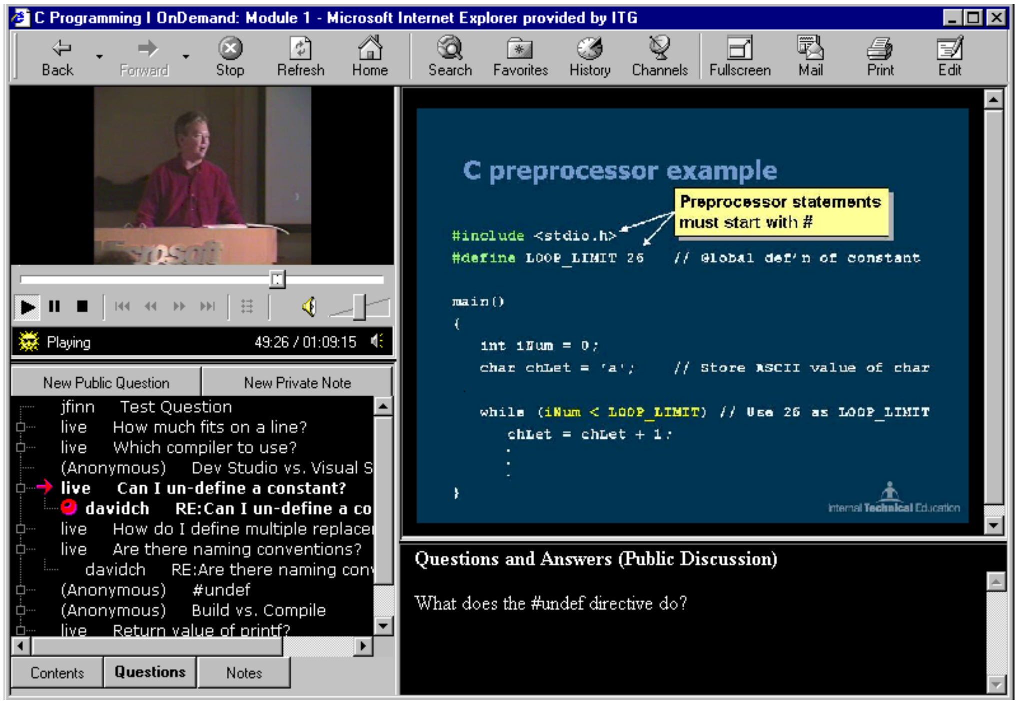

Figure 3: Web-based UI for On-Demand Education. This specialized display integrates course content and interactivity. Multimedia annotations support class discussion and a personal notebook, in addition to a table of contents and slides.

playing, and at the end of the demo video the main video should start playing again.

5.2 New UI Components

Based on these requirements, we designed and implemented a new set of interface components that shared the following properties:

The components are lightweight, self-contained, and completely web-based. In particular, we can embed multiple annotation displays in a single web page (for instance, in a frame set) and have each play a separate role.

The components can be configured with lightweight web page script (e.g., Javascript or VBScript). For example, we can specify which MRAS server to connect to and what annotations to retrieve when the web page containing the MRAS components loads.

The components support storing and displaying URL annotations. This is a particularly important annotation type, since it allows annotating video with anything that can be addressed by a URL and displayed (or executed) by a web browser. Slide switching and demo integration are accomplished using this annotation type.

The components expose the surrounding HTML Document Object Model (DOM) to contained annotations. By giving annotations access to the web page in which they are being displayed, the annotations can alter the behavior of the web page in interesting ways. This is how annotations containing demo videos automatically pause the main lecture video before starting the demo, for instance.

5.3 User Interface for On-Demand Education

Once implemented, we used these components, along with other standard web technologies, to compose a specialized web-based UI that met the user requirements for the on-demand education scenario. Based on informal user tests, we went through several iterations of the user interface before converging on the one shown in Figure 3. After describing the interface, we discuss alternative designs that we considered.

The lecture video at the top left is kept fairly small, as it is usually just a talking head. The frame in the top right is used for showing slides and/or demo videos. Slides are implemented as URL annotations (i.e., the appropriate segment of video is associated with the URL of the corresponding slide). Demo videos are implemented as annotations containing URLs to web pages which host the video content. When the main lecture video reaches a point where a demo should be played, it pauses and the demo is played in the slide frame. This frame has the largest area to enhance slide readability.

The lower left is devoted to three sets of annotations: A table of contents (labeled “Contents”), a shared question–and-answer set (labeled “Questions”), and personal notes (labeled “Notes”). All three sets occupy the same window, only one set shows at a time, and the user chooses which set is displayed by clicking on “tab” buttons at the base of the window. As the video plays, the closest annotation is highlighted with a red arrow. The contents of the highlighted annotation are shown in the preview pane in the bottom right. If the tab buttons are used to change the annotation set, the preview pane's content changes correspondingly. A user can also right-click on any annotation and seek to the corresponding point in video, reply to the annotation (thus creating a threaded discussion), edit it (if they have permissions), or delete it (if they are the owner). A single click on an annotation shows its contents in the preview pane and a double-click jumps the video to the point where the annotation was made.

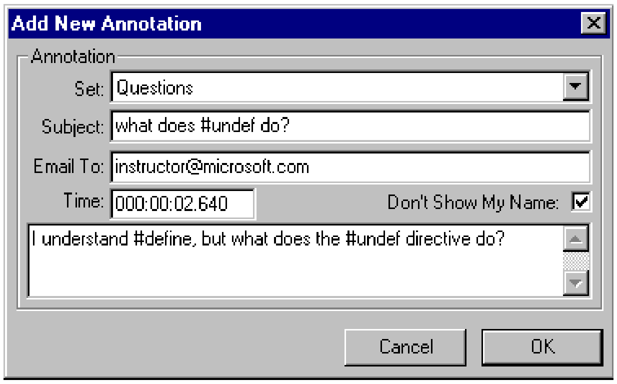

Figure 4: Dialog box for adding new annotations. Users can send annotations in email and create annotations anonymously.

Adding new annotations is initiated by clicking on one of the buttons just below the video frame. The left button is for adding to the shared discussion, and the right button is for creating private notes. In each case, a user is presented with a dialog box for composing the new annotation (Figure 4). Among other things, users can specify whether the annotation is to be anonymous and whether it is to be emailed to someone. There is no button for adding a new annotation to the table of contents since it is a read-only annotation set.

5.4 User Interface Design Tradeoffs

We reconsidered some aspects of the design based on the informal user tests we conducted:

We originally implemented an "add new annotation" input pane in the lower right hand corner of the UI (below the preview pane) to allowed users to type annotations naturally and continuously without having to open a separate dialog box each time. However, this unnecessarily consumed screen space and created serious mode problems. It was replaced by the add buttons below the video frame.

We were pushed toward simplicity over generality. To this end, we removed an option to add voice annotations, the ability to edit start- and end-points for annotations, and the ability to change which annotation set a new annotation should get added to. In general, we tried to reduce the gadgetry on behalf of the educational substance [13].

We experimented with different interaction models for common activities. For instance, originally, clicking an annotation once highlighted the annotation but caused no further action; and double-clicking it caused it to be opened in a separate read-only dialog box. We discovered that users preferred a single click to display the annotation in the preview pane, and a double-click to seek the presentation. This allowed users to browse through annotation contents without having the lecture video jump.

The original demo video annotations were too automatic for most users. That is, they were originally designed to automatically pause the main lecture video, load and play the demo in the slide frame, and afterwards automatically resume playing the main lecture video. This confused early test users, so we revised the behavior to pause the main video and load the demo, but to wait for a user to explicitly initiate demo playback.

6. On-Demand Training Study Goals

Our main goal was to compare the acceptability of asynchronous education and collaboration to that available in "live" classes. To this end, we studied a four-session introductory corporate training course in the C programming language, comparing one live and two on-demand instances of the course, and assessing their relative acceptability to students and instructors.

We did not contrast educational outcomes between the live and on-demand courses. Students come to corporate training courses with significantly different prior knowledge and educational objectives; to measure outcome accurately would require a much larger study and a change in the nature of the course. Also, as noted in the introduction, positive educational outcome has been well established: Hundreds of studies of broadcast and on-demand instruction using analog video overwhelmingly show that learning is comparable to -- or even better than in -- live classes [34]. The major remaining problem cited in these studies is dropout rate, which is likely related to the student attitudes that we measure in our study. Enhancing the student experience in on-demand courses with better support for collaboration around multimedia could address this challenge.

To assess whether student experience was enhanced by the interface for on-demand education we developed using MRAS, we wanted to answer the following questions:

- How convenient was the on-demand format? Did students really exploit it?
- Did the instructor save time because he did not have to teach a live class, or did answering online-questions take-up an equivalent amount of time?
- There is a fairly high attrition rate associated with corporate training classes. How did it compare between the two styles of offerings?
- Given the collaboration features provided by MRAS, was class participation comparable?

- Instructors often like to teach live classes because of interaction they have with students. How satisfied did they feel with the interaction arising in the on-demand class?
- What was the overall satisfaction of students with the on-demand course and collaboration features?

7. Study Procedure

To conduct our study, we observed and video-taped a "live" C Programming course conducted by the Microsoft Technical Education group (MSTE) and attended by employees. We used the video tapes, plus slides and other course content, to conduct two on-demand versions of the course using our system.

7.1 Live Course

The next "live" C Programming course is continually advertised to employees on MSTE's intranet website. Students enroll after obtaining their supervisor's permission. The course is taught in four two-hour sessions, held during normal business hours over a two week period. In the instance of the course we recorded, video cameras were placed at the back and front of the classroom to capture the instructor and the students. Students filled out a background questionnaire at the beginning of the course and a 12-question survey after each class session. At the end of the course, they filled out a 20-question survey. The instructor answered similar surveys to gauge his experience teaching the course.

7.2 On-Demand Course

The two on-demand courses we conducted were also advertised on the MSTE internal website. The first was also advertised on several internal email aliases. Subsequent "live" versions of the course were being offered at the same time as both of our on-demand versions, so students had a choice between "live" and on-demand when enrolling. Students were promised an MRAS T-shirt for participating in the on-demand courses.

Each lecture video from the live class was converted into a web-page with synchronized slides and tables of contents as shown in Figure 3. When the “Contents” (TOC) tab was selected, the preview pane showed the instructor’s detailed slide notes.

The shared discussion space was seeded with annotations containing questions that had been asked in the live class, to show students how annotations would look. All students had access to the shared discussion set, and each had a personal notes set. Annotations created in the shared discussion space during the first on-demand course were removed before the second course, so that students in each class started with the same seed annotations

Each on-demand course was taught over two weeks. It began with a live face-to-face session, during which we demonstrated MRAS, the students answered a background questionnaire, and the instructor gave a brief introduction to the course content. The course ended with another live session, during which we had the students fill out a 33-question survey.

During the course, students watched lectures from their desktop computers when they wanted, except that they were paced: They were asked to finish the first two sessions by the end of the first week, and the other two by the end of the second week. Halfway through the course, we asked students to fill out a 14-question web-based survey to gauge their progress. We had debated the pacing restrictions, but given the modest class size, we felt that if people’s viewing was too far spread apart, they would not benefit from each other’s comments. This may not be a issue in eventual large-scale deployments.

8. Results

Students found the on-demand format very convenient. 20 out of 21 students in the first on-demand course, and 11 out of 13 in the second, stated that time convenience had a large (positive) effect on their experience. This was also exhibited in the activity log: Students in the first and second on-demand courses watched an average of 65% (std. dev. = 0.32) and 72% (std. dev. = 0.32) of the course video, respectively, and used the navigational features to skip the parts they did not need to watch. Logons to the MRAS server per user per day (Figure 5) shows a relatively even distribution of connections throughout the courses, suggesting that students took advantage of the on-demand nature of course delivery. Peaks shown in Figure 5 at the beginning and end of the courses may illustrate the effect of enthusiasts and procrastinators.

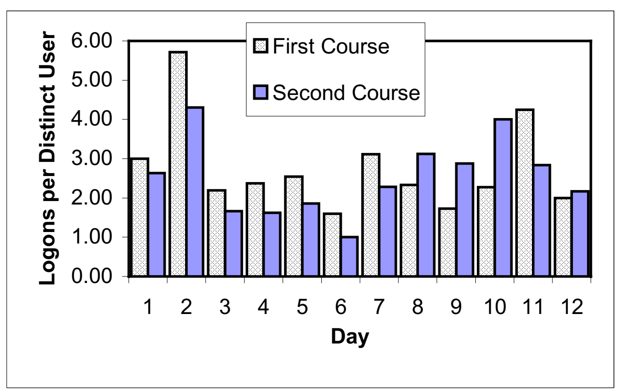

Figure 5: Logons per User per Day. Students took advantage of the on-demand nature of the course by “attending” class at their convenience.

Turning to instructor efficiency, in the live case the instructor spent 6.5 hours lecturing (which does not include preparation time and time spent commuting to and from the classroom), and there were no email questions.

For the on-demand classes we asked the instructors to closely monitor the time they spent handling student questions. They spent 1 hour each for the first and last live sessions, and in addition instructor 1 estimated 1 hour answering questions asked via annotations during the whole course, and instructor 2 estimated 2 hours. Both instructors felt that they answered student questions promptly and satisfactorily. All told, instructor 1 spent a total of only 3 hours teaching the on-demand course, and instructor 2 spent only 4 hours. The time savings could be even larger if face-to-face sessions were eliminated.

Student attrition rate (i.e., the ratio of people who started the courses but did not finish them) was lower in the on-demand courses. In the live course we observed, 19 out of 33 people, or about 58%, dropped out. In the on-demand courses, only 14 out or 35 (40%) dropped out of the first, and 7 out of 23 (39%) dropped out of the second. This is promising, but the result is tentative: Students chose the on-demand format over the alternative live format, so self-selection may have played a role.

Students in both on-demand courses felt they participated at roughly the same level as they had in past "live" course they took. This is partly supported by the data in Table 1, showing the number of content-related questions, procedural questions, comments, and answers for each of the courses. Although the average numbers for on-demand courses are smaller, the difference may be explained by the seeding of the on-demand lectures with questions from the live class. When we asked students in on-demand courses why they didn't ask more questions, the top two responses were that the material was clear, and that someone else had already asked the question they would have asked. If we add the live and on-demand annotations (right two columns in Table 1) the total level of interaction in the on-demand classes is higher than in the live class. Taking a long-term perspective, the best questions from a whole series of class offerings could be accumulated in an annotation database, so that the experience of an on-demand student is significantly better than that of live students.

Table 1: Comparison of content-questions, procedural-questions, comments, and answers between courses. "O.D" means on-demand. 'per-student' statistic was calculated by dividing TOTAL by the number of students who finished the course.

	Live	O.D. 1	O.D. 2	O.D. 1 + Live	O.D. 2 + Live
Content	15	5	5	20	20
Procedural	0	2	2	2	2
Comments	0	4	2	4	2
Answers	15	17	9	32	24
TOTAL	30	28	18	58	48
per student	2.14	1.33	1.29	2.76	3.43

As for the *value* of class participation, overall students in each course reported no difference in the quality of interaction. Based on the assumption that students who are more familiar with course content are better judges of the value of class participation, if we consider only those students who knew 20% or more of the content before taking the course (which was 57% of the "live" class, and 76% and 50% of the on-demand students), on-demand students valued other students' comments significantly more (using one-way analysis of variance, ANOVA, on survey answers, we found p=0.014) than students in the live class did. These numbers are presented as part of Table 2. One student liked seeing others' input because "[he] learned something [he] didn't even think of," while another said the student comments "better explained the issue at hand." Another student remarked that the collaborative features "...helped me compare myself to others in the group. Sometimes I'd ask myself something [and it] was nice to see I had the right answers."

Table 2: Survey Results. Probability p was calculated using one-way analysis of variance (ANOVA). Items marked with * were calculated for students who knew more than 20% of material before the course began (the means are across all students though). "O.D" means on-demand.

Category		Live	O.D. 1	O.D. 2	p
Pace 1=very slow, 5=very fast		3.19	2.90	3.04	n/a
Paying Attention	% Close	67.50	59.05	61.92	n/a
	% Moderate	23.79	26.90	28.46	n/a
	% Not	8.71	14.05	9.62	n/a
How much learned? 1=much less than usual, 5=much more than usual		2.83	3.65	3.5	0.033
Satisfaction with... 1=v. dissatisfied, 5=v. satisfied	Quality	3.82	4.14	4.15	0.055*
	Content	3.64	3.86	4.31	0.007*
	Time	3.89	4.35	4.08	0.016*
Value of other students' comments 1=definitely not valuable, 5=definitely valuable		3.00	3.38	3.35	0.014*
presentation format interfered with ability to learn 1=strongly interfered, 5=strongly enhanced		2.07	3.71	3.54	0.000
Instructor was accessible and responsive 1=strongly disagree, 5=strongly agree		4.29	3.43	3.31	0.002

When we assessed instructor satisfaction with the on-demand format, they responded that they had too little contact with students and did not get enough feedback to know how well students were doing. On the other hand, the instructors liked the convenience and efficiency of the on-demand course format.

Students in the on-demand courses reported significantly lower instructor responsiveness than students in the live class. However, they also reported liking the presentation format of the course significantly more.

When we asked students in all courses whether they were satisfied with lecture quality, course content, and use of time, there was no difference between on-demand and live student responses (Table 2). When we again limited the student pool to those who knew more than 20% of the course content before starting the course, however, we found that on-demand students appreciated these things more than students in the live course.

9. General Feedback

We received several interesting comments during the final 1-hour sessions of the on-demand courses:

Students indicated that the value of the on-demand course would be significantly enhanced if they could have participated from home (we used 110Kbps audio-video, so modem users could not access it). They were willing to put up with audio-only for that flexibility.

A majority of students took personal notes on hard copies of the course workbook, instead of using MRAS. Key reasons for doing this included 1) no guarantee that MRAS notes would be available in the future; 2) The convenience of paper; and 3) No easy way to print the notes they took with MRAS.

Students would have liked to be able to annotate slides and workbook content, and not just link annotations with the timeline of the video. Creating a system and interface for fully general annotation of mixed-media documents is an important direction for future work.

Students liked asynchrony, but they missed 1) the immediate answers to questions in live classes, and 2) the back-and-forth of interactive exchange. To address the first concern, they suggested posting questions to an email alias or newsgroup, so that a group of TAs or “knowledgeable” fellow students could provide instantaneous replies. To address the second concern, they suggested online office hours, where people could participate in interactive chat.

Instructor comments were more limited. A key concern was to increase interaction with the students. One instructor said that to some extent he felt like a glorified grader or TA, which is not as rewarding. This is a genuine concern, as instructors are gatekeepers to the wider adoption of this kind of technology.

10. Concluding Remarks

There is a growing interest in scaling educational systems to reach large numbers of students in a cost-effective manner and without negatively impacting learning. This scaling will more likely come via asynchronous on-demand systems than through synchronous “live” systems. Previous work has shown that a key challenge for the asynchronous educational model is to maintain student engagement: when students continue to participate, their learning is the same as with classroom attendance. Promoting collaboration around on-demand multimedia educational content can address this challenge.

We have investigated how a system that couples multimedia annotations with web technologies and email can support rich interaction in asynchronous environments. We discussed the requirements and user interface design for on-demand videotaped lectures and the results from a first case study. Key elements included scriptable web-based components that can be embedded in browser frames and can implicitly connect to an annotation server without user involvement. Once we had built these components, an analysis of user requirements led to the design and implementation of an interface to support on-demand training. The interface performed reasonably well in our study and according to participants met the basic requirements.

The case study showed a positive outcome for most of the issues we explored. Students felt they benefited from the on-demand delivery, instructors spent less time and were more focused on questions, the attrition rate was lower than that for live classes, and the participation level was felt to be comparable to live courses previously taken by the students. One student said “I would definitely take another MRAS course, it was great and easy to use.” Another said, “I really enjoyed this! Thank you so much for doing [C Programming I]! Now if only [C Programming II] were available…:-)”. Yet another said “This was a fantastic course. Everyone I've mentioned it to, or showed it to, thinks it is awesome and would increase the [number of] classes they attend!” However, the instructor concerns remain.

The current system represents a starting point. By learning from ongoing use, we should be able to significantly enhance student experience with multimedia in the on-demand education and training arena.

11. Acknowledgements

Thanks to David Aster, Barry Preppernau, Catherine Davis, Carmen Sarro, Shanna Brown, David Chinn, and Steven Lewis, all from Microsoft Technical Education, for their help in conducting the C Programming Language courses. Thanks to Suze Woolf and Jonathan Cluts for their UI improvement suggestions.

12. References

[1] Abowd, G., Atkeson, C.G., Feinstein, A., Hmelo, C., Kooper, R., Long, S., Sawhney, N., and Tani, M. Teaching and Learning as Multimedia Authoring: The Classroom 2000 Project, Proceedings of Multimedia ’96 (Boston, MA, Nov 1996), ACM Press, 187-198.

[2] Ackerman, M.S., and Malone, T.W. Answer Garden: A Tool for Growing Organizational Memory, Proc. ACM Conference on Office Information Systems, 1990, 31-39.

[3] Bargeron, D., Gupta, A., Grudin, J., and Sanocki, E. Annotations for streaming video on the Web: system design

and usage studies. Proceedings of the Eight International World Wide Web Conference (Toronto, May 1999).
[4] Bates, A.W. The Future of Education, Telektronikk, Vol 92, No 3-4, 1996. pp 82-91.
[5] Carrer, M., Ligresti, L., Ahanger, G., and Little, T.D.C. An Annotation Engine for Supporting Video Database Population, Multimedia Tools and Applications 5 (1997), Kluwer Academic Publishers, 233-258.
[6] Computer Supported Intentional Learning Environments (CSILE) project. http://csile.oise.on.ca/.
[7] Cordero, C., Harris, D., and Hsieh, J. High speed network for delivery of education-on-demand, Proceedings of the International Society for Optical Engineering Conference (San Jose, CA, January 1996), 161-72.
[8] Davis, and Huttonlocker. CoNote System Overview. (1995) Available at http:// www.cs.cornell.edu/ home/ dph/ annotation/ annotations.html.
[9] Hypernix Gooey web communication system. http://www.hypernix.com/.
[10] Kim, K.W., Kim, K.B., Kim, H.J. VIRON: An Annotation-Based Video Information Retrieval System, Proceedings of COMPSAC '96 (Seoul, Aug. 1996), IEEE Press, 298-303.
[11] Laliberte, D., and Braverman, A. A Protocol for Scalable Group and Public Annotations. 1997 NCSA Technical Proposal, available at http:// union.ncsa.uiuc.edu/ ~liberte/ www/ scalable-annotations.html.
[12] Lawton, D.T., and Smith, I.E., The Knowledge Weasel Hypermedia Annotation System, Proceedings of HyperText '93 (Nov 1993), ACM Press, 106-117.
[13] LeDuc, A.L. What is truly important about information technology in higher education, Journal of Computing in Higher Education, Vol 8, No 1, 1996. pp. 124-39.
[14] Lee, S.Y., and Kao, H.M. Video Indexing – An Approach based on Moving Object and Track, Proceedings of the SPIE, vol. 1908 (1993), 25-36.
[15] Lotus Learning Spaces. http://www.lotus.com/home.nsf/welcome/learnspace.
[16] Marsh, L.C., and Wells, K.L. Key Aspects of a Computerized Statistics Course, Journal of Computing in Higher Education, Vol 8, No. 2, 1997. pp.72-93.
[17] Marshall, C.C. Toward an ecology of hypertext annotation, Proceedings of HyperText '98 (Pittsburgh, PA, USA, June 1998), ACM Press, 40-48.
[18] Maurer, H., Lennon, J. Digital libraries as learning and teaching support, Proceedings of the International Conference on Computers in Education (Singapore, December 1995), Association for the Advancement of Computer Education, 28-33.
[19] McKendree, J., Stenning, K., Mayes, T., Lee, J., and Cox, R. Why Observing a Dialogue may Benefit Learning, Journal of Computer Assisted Learning, 14, 1998, 110-119.
[20] NovaWiz Odigo web communication system. http://www.odigo.com/.
[21] Organik FAQ builder. http://orbital.co.uk/organik_info/index.htm.
[22] Parker, A. Distance Education Attrition, International Journal of Educational Telecommunications, Vol 1, No 4, 1995. pp. 389-406.
[23] Roscheisen, M., Mogensen, C., Winograd, T. Shared Web Annotations as a Platform for Third-Party Value-Added, Information Providers: Architecture, Protocols, and Usage Examples, Technical Report CSDTR/DLTR (1997), Stanford University. Available at http://www-diglib.stanford.edu/rmr/TR/TR.html.
[24] Scardamalia, M., & Bereiter, C. Student communities for the advancement of knowledge. Communications of the ACM, Vol. 39, No. 4, 1996. pp. 36-7.
[25] Schickler, M.A., Mazer, M.S., and Brooks, C., Pan-Browser Support for Annotations and Other Meta Information on the World Wide Web. Proceedings of the Fifth International World Wide Web Conference (Paris, May 1996), available at http:// www5conf.inria.fr/ fich_html/ papers/ P15/ Overview.html.
[26] Smith, B.K., and Reiser, B.J. What Should a Wildebeest Say? Interactive Nature Films for High School Classrooms, Proceedings of ACM Multimedia '97 (Seattle, WA, USA, Nov. 1997), ACM Press, 193-201.
[27] Stanford Online distance education service. http://stanford-online.stanford.edu/.
[28] ThirdVoice web annotations. http://www.thirdvoice.com/.
[29] uTok web annotation system. http://www.utok.com/.
[30] Vetter, R.J., and Jonalagada, C. A multimedia system for asynchronous collaboration using the Multicast Backbone and the World Wide Web, Proceedings of the First Annual Conference on Emerging Technologies and Applications in Communications (Portland, OR, May 1996), IEEE Computer Society Press, 60-63.
[31] Vicarious Learning Project. http://www.hcrc.ed.ac.uk/gal/vicar/.
[32] WebCT e-Learning Hub. http://www.webct.com/.
[33] Weber, K., and Poon, A. Marquee: A Tool for Real-Time Video Logging, Proceedings of CHI '94 (Boston, MA, USA, April 1994), ACM Press, 58-64.
[34] Wetzel, C.D., Radtke, P.H., & Stern, H.W. Instructional effectiveness of video media. (1994). Erlbaum.
[35] Zadu web customer contact system. http://www.zadu.com/.