

# Asynchronous Collaboration Around Multimedia Applied to On-Demand Education

David Bargeron, Jonathan Grudin, Anoop Gupta, Elizabeth Sanocki, Francis Li<sup>1</sup>, Scott  
LeeTiernan<sup>2</sup>

Collaboration and Multimedia Systems Group  
Microsoft Research, Redmond, WA 98052

{davemb, jgrudin, anoop}@microsoft.com, lizsanocki@hotmail.com, fli@cs.berkeley.edu,  
slt@u.washington.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 potentially a significant tool for delivering online educational content effectively. A multimedia annotation system tightly integrated with email provides a powerful platform on which to base such functionality. In this paper we describe a series of studies of such a system. First, we built a prototype annotation system and refined it based on results of laboratory tests. We then extended the system to support asynchronous collaboration for on-demand training and studied its effectiveness in two corporate training courses, assessing student experience, instructor experience, and user interface appropriateness. Having identified possibilities for enhancing engagement and collaboration with the tool, we conducted another set of laboratory studies. Through this iterative process we are creating a platform and identifying processes for its use that enable students and instructors to exploit the advantages of asynchronous education while compensating for the reduction in face to face interaction.*

## Keywords

Annotation, Multimedia, Asynchronous Collaboration, Distance Education, Workplace Training, Group Interaction.

## 1. Introduction

Many see the Web as an irresistible platform for education and workplace training [5][21], 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 are broadcast via TV to remote students' sites and telephone back-channels were provided for questions to the instructor. Hundreds of studies have shown that students can learn as much from such broadcast lectures as from live classroom attendance [40]. 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

---

<sup>1</sup> Current address: Computer Science Department, University of California Berkeley, Berkeley, California

<sup>2</sup> Current address: Psychology Department, University of Washington, Seattle, Washington

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 [40]. But this relatively low-tech approach suffers from decreased engagement: 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., [33]) 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 the evolution of MRAS (*Microsoft Research Annotation System*), our general architecture for multimedia annotations, focusing on extensions to address short-comings of multimedia use in on-demand education. We built a set of closely integrated, independently reusable client components that serve as a flexible platform for asynchronous collaboration. The programmable components can be embedded in web pages.

We have studied, refined, and extended MRAS over a two-year period. First we carried out two laboratory studies to refine the functionality and interface of the system. We then consulted with instructors and students of corporate training courses to gather requirements for its use in that setting. We used the MRAS components to produce an application and interface that were then examined in two on-demand offerings of a popular introductory training course. Based on the outcome of these trials, we designed features to promote engagement through collaboration in group exercises and embedded self-assessment and returned to the laboratory to examine possible processes for using them.

In this paper we emphasize the use of MRAS in regularly scheduled corporate education courses. 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 briefly review the initial laboratory tests of MRAS and describe the extensions intended to support on-demand training. We describe the study design and findings. Finally, we briefly describe the subsequent study of features added to enhance engagement and assessment, and discuss future directions.

## 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 [12] and Animal Landlord [32] 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 [22] 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 [33][36][9]. Although MRAS is a powerful system for storing indices, our focus is on more dynamic, asynchronous content.

WebCT [38] and Lotus LearningSpace [20] 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). The CoWeb projects at Georgia Tech [10] explore similar functionality. None offer MRAS’s rich support for multimedia, for video annotations, and for fine-grained organization and sharing of annotations.

AnswerGarden [2] and Organik [27] 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 [23] 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 [37] explored several ideas related to our work and demonstrated that student dialogs can be a valuable learning resource for subsequent students [24]. The Computer Supported Intentional Learning Environments (CSILE) project [8] and its commercial outgrowth called Knowledge Forum have shown that collaborative learning can enhance the depth of each students’ understanding of the subject matter [30]. 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 [31] and NCSA [15] have proposed scalable Web-based architectures for sharing annotations on web pages. CritLink [11] is a web note-sharing system; ThirdVoice [34], NovaWiz [26], Hypernix [13], uTok [35], and Zadu [41] have all recently released commercial systems. These are similar in principal to MRAS, but do not support fine-grained access control, annotation grouping, video annotations, and rich annotation positioning. The web-based Knowledge Weasel [16] 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 [29] 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 [18], Marquee [39], VIRON [14], and VANE [6]. 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.

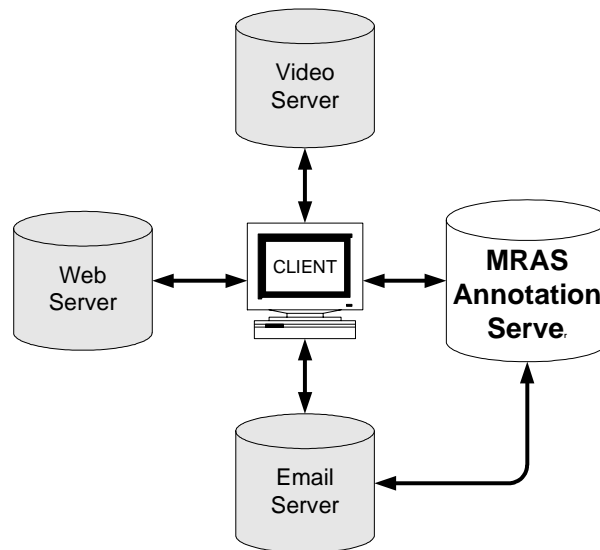
### **4.1 Multimedia Annotations**

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 MRAS Annotation Server 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 MRAS User Interfaces

Figure 2 shows the versions of the MRAS interface used by students in the corporate training classes. Unlike our initial interface, it is entirely web-based: The display is a single browser window with several frames.

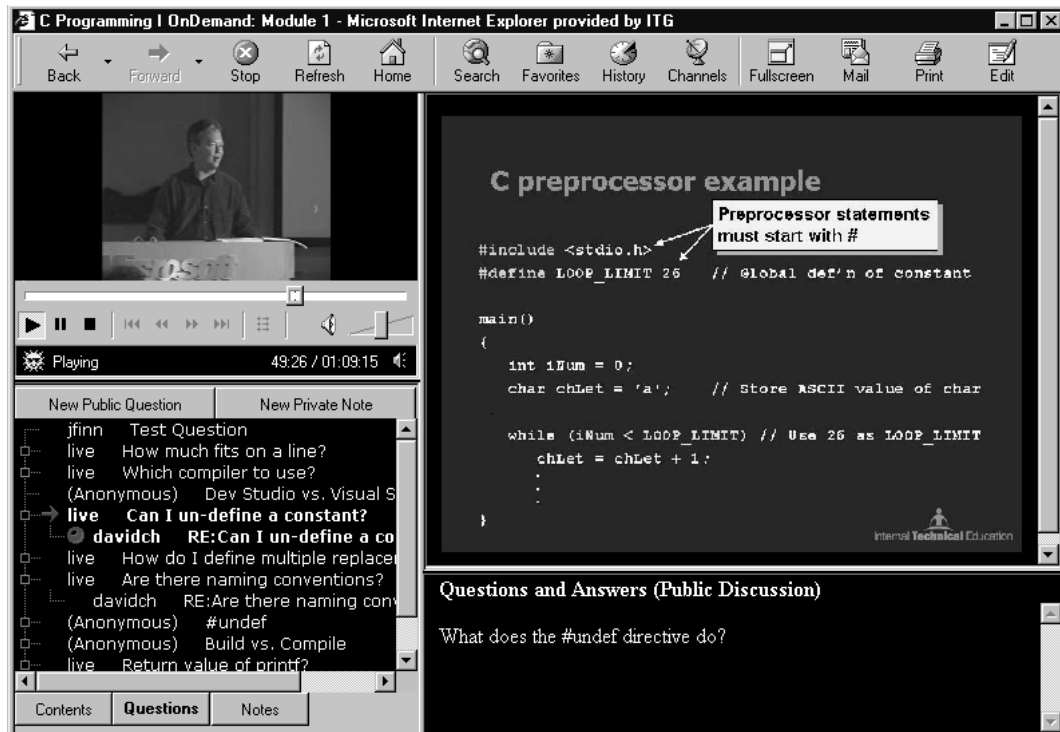


Figure 2: Web-based MRAS interface.

#### 4.4 Annotation System Features and Functions

The video in the upper left of the browser window in Figure 2 is displayed with a standard media player. The slides on the right are synchronized with the video. In the lower left is an annotation set: comments, questions, and replies made by previous viewers of the video. Each annotation is linked to the specific point in the video at which the video paused when the annotation was created. The red arrow points to the annotation that was created closest to the current position of the video; the reply below it is also highlighted. To the lower right is the annotation preview window, which automatically shows text associated with the currently highlighted annotation.

As the video plays, annotation titles scroll by. A viewer can click on an annotation and read it in the preview frame, right-click to seek to the point in the video where an annotation was created, or reply to it, forming a threaded discussion.

Currently displayed is the 'Questions' annotation set, which can be written to and read by all students in the class, in this case a lecture on transaction servers. The 'Contents' button will bring up a read-only set of annotations that consist of the slide titles, a table of contents for the lecture. Using these, students can skip from one topic to another. The 'Notes' button on the right brings up a set of personal notes for this particular viewer. Each private note, like other annotations, is linked to the point in the video at which it was created, much as a written note in a book is clearly linked to its context, the page or paragraph on which it is written.

New annotations can be added to a user's personal 'Notes' annotation set by clicking on the 'New Private Note' button. A new annotation can be added to the public 'Questions' annotation set by clicking on the 'New Public Question' button. The dialog box that appears for adding a new public question annotation is shown in Figure 3.

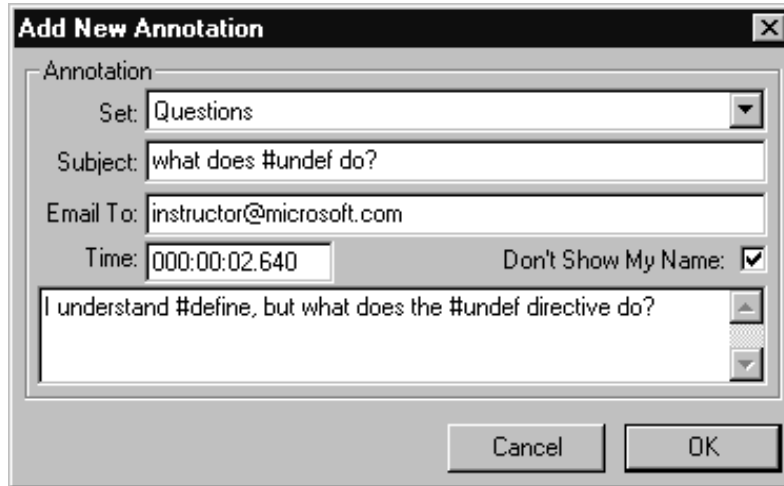


Figure 3: Dialogue box for adding a text annotation.

## 5. Lab Studies of Handwritten, Typed, and Voice Annotations

Our first studies were conducted to determine the potential for on-line in-context annotation of streaming video and to see how people used and responded to a range of features. These studies are sketched here and reported in detail in [3]. The MRAS interface was similar to that described above, but was not embedded in a single browser window; the annotation set opened in a separate window that floated above the web-based video. Boxes to enter URLs for the video and for the annotation set appeared in the interface. No slides accompanied the video, so there was no window or frame for slides. This interface appears in the Appendix (see Figure 7). One significant difference is that the interface supported spoken voice annotations. The dialogue box for annotations used in these experiments is shown in Figure 4; the voice annotation radio button has been selected in the lower left corner

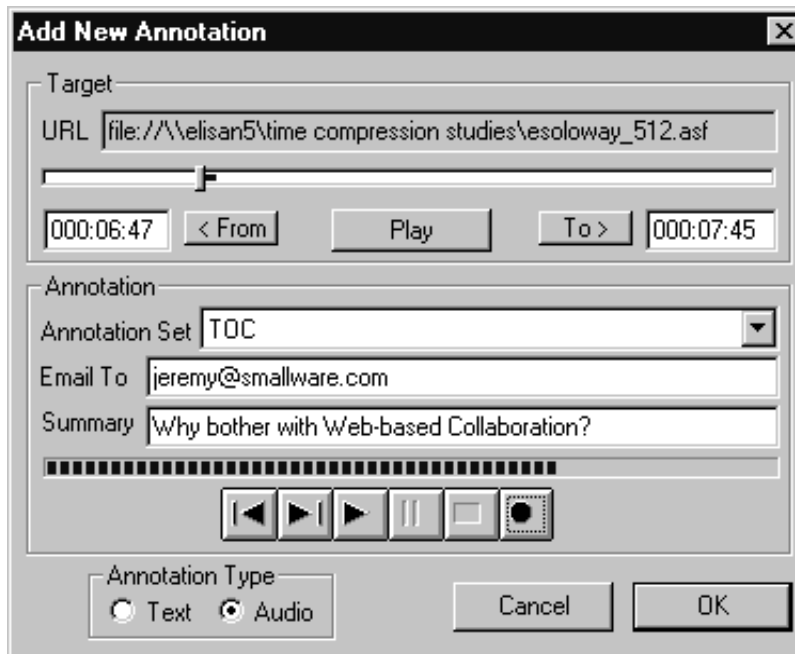


Figure 4: Dialogue box for adding a voice annotation in the first lab study.

The first of two studies contrasted handwritten and typed annotations; the second contrasted the use of text and audio annotations, and explored the creation of annotation threads through replies. These studies are reported in more detail in [3].

## 5.1 Participants

Six people participated in the first study, 18 in the second. All were intermediate to advanced Windows users with no involvement with the research. They received software products for their participation.

## 5.2 Method

Participants first completed a background questionnaire. They were given a brief introduction to the MRAS system before using it. During the experimental sessions behaviors such as pausing or rewinding the video were identified through a one-way mirror. Task time and annotation number and content were also recorded. After the study, participants completed a post-study questionnaire and were debriefed.

Participants were asked to assume that they were students in a course. To prepare for a discussion in the next class meeting, they needed to watch a video presentation of Elliot Soloway's ACM '97 conference talk "Educating the Barney Generation's Grandchildren." Their goal was to generate questions or comments from the video for the class discussion. Each was given as much time as needed to complete the task.

### 5.2.1 *Handwriting vs. online annotation*

In the first study, six participants made annotations in two ways: using handwritten notes and using text-based MRAS notes. Half took handwritten notes during the first half of the video and switched to MRAS for the second half of the video, half used MRAS for text-based note taking during the first half, and were instructed to take pen and paper notes as they normally would for the second half.

### 5.2.2 *Text vs. audio in threaded discussions*

In the second study, 18 participants used MRAS to take notes for the entire video and were told that they were adding their comments and questions to a shared set of notes. Each participated in one of three conditions, a Text-Only condition where they only added text annotations, an Audio-Only condition, and a Text-and-Audio condition, where both text and audio annotations were allowed.

For each condition, six participants participated in sequence with their notes saved to a common annotation set, so that each subsequent participant in that condition could see, review, and respond to the annotations created by previous participants. The first participant in each condition started with a set of "seed" annotations created by a participant in an earlier pilot study and adapted to the appropriate annotation medium for each condition. Participants were asked to review existing annotations before adding their own to avoid redundant annotations.

## 5.3 Results

### 5.3.1 *Handwriting vs. online annotation*

Details of this data analysis are published in [3]. We review the findings to establish the context for the subsequent studies.

Participants using MRAS took an average of 16.8 notes; participants using pen and paper took a statistically indistinguishable average of 16.2 notes. Taking notes using MRAS took a significant 32% longer than on paper on average. With paper, users often took notes while the video continued to play, thus overlapping listening and note-taking. For MRAS, the system pauses the video when adding (this feature is necessary for audio annotations and was used for text as well). However, this difference may not negatively affect the perception of MRAS---as noted below, all participants reported that the benefits of taking notes with MRAS outweighed the costs.

Order effects help explain these results. Participants who used MRAS first paused the video while taking paper notes; *none* of the paper-first participants paused. This suggests that prior to using MRAS people took notes as though in a live lecture; MRAS use reminded them of the ability to pause and not miss some of the lecture. As a consequence, the saving in time for paper notes was due to the paper-first participants. Another order effect is that MRAS-first participants made significantly more notes than those who first used paper: Experience taking paper notes seemed to reduce the inclination to take notes using the system.

All participants expressed preference for MRAS to paper. Comments emphasized organization, readability, and contextualization. One participant stated that the notes she took with MRAS were “...much more legible, and easier to access at a later time to see exactly what the notes were in reference to.” Another said that MRAS “...allows me to jump to areas in the presentation pertaining to my comments.”

When asked which method resulted in notes that would be more useful 6 months down the road (for instance for studying for an exam), all participants again chose MRAS, citing better organization, increased readability, and automatic positioning of notes within the lecture video. In this regard a participant said of her notes “the end product is more organized...The outline feature...allows for a quick review, with the option to view detail when time permits.” Participants noted that the seek and tracking features of the MRAS user interface were particularly useful for relating their notes to the lecture video. Although the participant population was small, these results encouraged us to conduct more extensive laboratory tests and field studies described below.

### 5.3.2 Text vs. audio in threaded discussions

Previous studies have compared the use of text and audio for annotation of text [7][25], but not annotation of Web-based video.

There was a fairly high rate of participation in all three conditions: 50 annotations were added in Text-Only, 41 in Audio-Only, and 76 in the Text and Audio condition. These differences, and length of time to complete the exercise, were not significantly different across conditions. (For more detail see [4].)

Participants in the Text and Audio condition used text more than audio for new annotations and replies. They were more likely to use text to reply to text and equally likely to use text in response to an audio annotation. When asked which medium they found easier, 4 out of 6 chose text, despite the relative ease of speaking over typing. One said that typing text “...gives me time to think of what I want to put down. With the audio, I could have paused, gathered my thoughts, then spoke. But, it is easier to erase a couple of letters... than to try to figure out where in the audio to erase, especially if I talked for a while.”

Participants in the Text and Audio condition were much more likely to reply to text annotations in the first place. User feedback from both the Text and Audio and Audio-Only conditions explains why: Participants generally felt it took more effort to listen to audio than to read text. One participant in the Text and Audio condition was frustrated with the speed of audio annotations, saying that “I read much faster than I or others talk. I wanted to expedite a few of the slow talkers.” Another participant pointed out that “it was easy to read the text [in the preview pane] as the video was running. By the time I stopped to listen to the audio (annotations) I lost the flow of information from the video.”

As a result of this study, we have restricted our focus to text annotations in subsequent studies. Should automated speech-to-text translation become practical, voice annotation will be worth revisiting.

With participants participating sequentially, we expected to see the number of new annotations drop off as the study progressed, with earlier participants annotating the most salient parts of the video, leaving less for later participants to add. In fact, the number of annotations increased significantly, due to an increase in the number of replies; more interaction occurred as the study progressed (Figure 5).




Figure 5: Annotations Added Over Time. A linear least-squares trend line is superimposed in black.

This curve must asymptote or even decline as the number of previous viewers continues to increase, but it is an open, significant research question as to how viewers will respond to large numbers of annotations. In a study of email threads, it was found that the likelihood of another contribution increases with the number of contributions up to about 6. With MRAS, there is both an overall set of annotations and threaded responses to specific annotations. In a class situation, other factors could influence the likelihood of contributing a public annotation. The later one views material, the fewer the number of students who will encounter a comment later, and procrastinators viewing the material very late may not have the time to read and reply to annotations. We return to this topic in the final study concluding this paper.

As in the first study, users generally positioned an annotation 10 to 15 seconds after the relevant part of the video because of the time to decide to respond and press the “Add” button. Contrary to our expectation, most subsequent participants were not distracted by this. When watching the lecture video from beginning to end, they found it natural to see an annotation display in the preview window several seconds after the relevant part in the video.

Participants generally liked using MRAS for taking shared notes. Asked whether they made more comments and questions using MRAS than they would have in a live lecture, 14 of 18 said yes. Of course, this was due in part to the fact that they were watching a video that could be paused and replayed.

Asked whether prior annotations helped or distracted them, some reported that the additional information was confusing because it was not perfectly synchronized with the activity in the lecture video. However the majority, across conditions, found others’ annotations useful in guiding their own thinking, as reported for annotations of text [22]. One said that it was interesting and useful to see “...what other people find important enough to comment on.” Others: “It was thought-provoking to read the comments of others.” “It was more of an interactive experience than a normal lecture. I found myself reading the other notes and assimilating the lecture and comments together.” “...gives you the opportunity to see what others have thought, thus giving you the chance to prevent repetition or [to] thoughtfully add on to other’s comments.”

Finally, participants were asked whether they found using MRAS to be an effective way of preparing for a class discussion. A strong majority across all conditions agreed. The Text and Audio participants averaged 6.0 where 7 was “strongly agree” and 1 “strongly disagree.” The average in the other conditions was 5.0 out of 7. One was particularly enthusiastic: “I think that this software could really get more students involved, to see different viewpoints before classroom discussions, and could increase students’ participation who would otherwise not get too involved.” The research that followed was undertaken to see if this observation might be correct.

## 6. MRAS use in Corporate Training Courses

Our next step was to deploy MRAS in multi-session corporate training courses. This study is drawn from the account in [4]. To design the new features and interface (Figure 2), we drew on the studies just described, and on our observations of courses and discussions with instructors.

### 6.1 User and Task Requirements

- **Simplify the annotation user interface:** The original interface required too many decisions from users, such as which annotation server to connect to, which annotation sets to retrieve, and which annotation set to add new annotations to. (For details see Appendix.) We concluded that an author of a class web page should make such decisions and present defaults to students. For the classes in the following study, students are restricted to three annotation sets (read-only table of contents, personal notes, class discussion), all accessible with a button click. Access to functions such as adding annotations is more visible. Only text annotations are supported; as noted, they were not used heavily in the lab study, and in broad use could not be heard on speakerless systems.
- **Integrate annotation features directly into the on-demand education interface:** The original interface did 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 (as in Figure 7). Users are sometimes confused by this interference. We extended the design so that interface components are embedded within a regular browser frame.
- **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 2), 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 playing, and at the end of the demo video the main video should start playing again.

## 6.2 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.

In addition to supporting text (and, in theory, voice) annotation, the components support storing and displaying URLs as annotations. This 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. However, this is accomplished outside the students’ view; they simply see slides and demos appear when appropriate, and did not have access to voice annotation.

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.

## 6.3 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 requirements for the on-demand education scenario. Based on informal user tests, we went through several iterations before converging on the one shown in Figure 2. 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.

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 3). 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.

## 6.4 User Interface Design Tradeoffs

We reconsidered some aspects of the design based on our informal tests:

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 [17].

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.

## 6.5 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 [40]. 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?

## 6.6 Study Procedure

To conduct our study, we observed and video-taped a "live" C Programming course conducted by the Microsoft Technical Education group and attended by employees. The course is taught in four two-hour sessions, held during normal business hours over a two week period. 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. We used the video tapes, slides and other course content to conduct two on-demand versions of the course using our system.

The on-demand courses were listed on the same internal website used to enroll students to live courses. The first was also promoted on several internal email aliases. Subsequent live versions of the course were offered at the same time as each of our on-demand versions, so students had a choice 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 2. 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 debated the pacing restrictions, but given the modest class size, we felt that if people's viewing was too far spread apart, they would benefit less from each other's comments. This may be less of an issue in eventual large-scale deployments, although the final study looks at student exercises, a more natural way to push students to engage with lecture material before the last minute.

## 6.7 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 6) 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 6 at the beginning and end of the courses may illustrate the effect of enthusiasts and procrastinators.

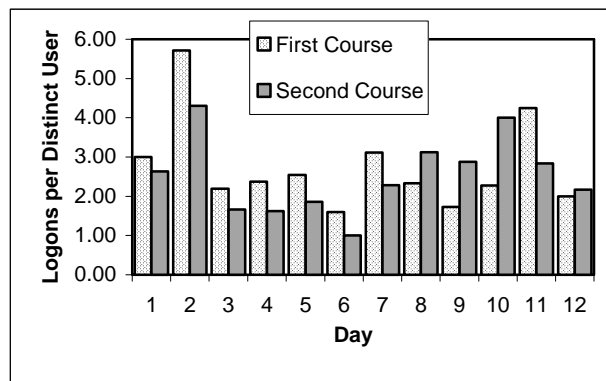


Figure 6: Logons per user per day.

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 of 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.

|                    | Live | O.D. 1 | O.D. 2 | O.D. 1 + Live | O.D. 2 + Live |
|--------------------|------|--------|--------|---------------|---------------|
| <b>Content</b>     | 15   | 5      | 5      | 20            | 20            |
| <b>Procedural</b>  | 0    | 2      | 2      | 2             | 2             |
| <b>Comments</b>    | 0    | 4      | 2      | 4             | 2             |
| <b>Answers</b>     | 15   | 17     | 9      | 32            | 24            |
| <b>Total</b>       | 30   | 28     | 18     | 58            | 48            |
| <b>Per student</b> | 2.14 | 1.33   | 1.29   | 2.76          | 3.43          |

Table 1: Comparison of live and On Demand (O.D.) content and procedural questions, comments, and answers. 'Per-student' calculated by dividing Total by number of students finishing the course.

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."

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.

| Category  |            | Live  | O.D. 1 | O.D. 2 | p      |
|---|------------|-------|--------|--------|--------|
| Pace<br>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?<br>1=much less than usual,<br>5=much more than usual                                |            | 2.83  | 3.65   | 3.5    | 0.033  |
| Satisfaction with...<br>1=v. dissatisfied,<br>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<br>1=definitely not valuable,<br>5=definitely valuable              |            | 3.00  | 3.38   | 3.35   | 0.014* |
| presentation format interfered with ability to learn<br>1=strongly interfered,<br>5=strongly enhanced |            | 2.07  | 3.71   | 3.54   | 0.000  |
| Instructor was accessible and responsive<br>1=strongly disagree,<br>5=strongly agree                  |            | 4.29  | 3.43   | 3.31   | 0.002  |

Table 2: Survey Results. Probability based on one-way ANOVA. Items marked \* calculated for students who knew more than 20% of material before course (means are across all students).

## 6.8 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.

## 7. Fostering Engagement and Interaction in On-Demand Courses

A programming language course is more likely to draw student questions than it is to spur a threaded discussion; how might the technology fare with other kinds of material? We noted that the flexibility of being able to watch video on demand can lead some students to procrastinate. If students watch a lecture at the last minute they may not have the time to pay attention to comments left by other students, and it will be too late for other students to see any comments they make. In addition, although students used the system to take notes and ask questions, there may be less interaction among students when they do not meet face to face.

The final study reported here involves a return to the laboratory to explore MRAS enhancements intended to address these possible problems, drawing on a more detailed report in [19]. Specifically, we looked at the feasibility of using group exercises that an instructor could assign as a way of pushing students to view the material before the last minute and to increase interaction among class members.

In a real distance-learning course, groups of students assigned to do an exercise together might arrange to meet in real time (face-to-face or online), even when the entire class cannot meet; alternatively, group members may have to work exclusively asynchronously. To simulate these circumstances, we designed two group exercise conditions, one

in which participants met together live and used MRAS to view lectures and report results, the other in which they used MRAS to conduct the exercise asynchronously. A control condition in which participants completed the exercise alone was included. We also selected lecture material that is more conducive to discussion. Finally, we adapted MRAS by adding a ‘Group’ annotation set that can be viewed and added to by group members. This required buttons for viewing a ‘Group’ set and for adding a ‘New Group Comment.’

## 7.1 Procedure

Sixteen beginning to advanced computer users participated in this study in exchange for a software gratuity. Each was randomly assigned to the asynchronous group (N=6), “live” group (N=4), or “solo” (N=6) condition.

Participants role-played taking an economics course through a university distance-learning program. They were told they would use a system to view and annotate two short pre-recorded lecture videos and then complete an assignment. Each completed a short training session in which they learned the core functionality of MRAS.

Participants then watched two 8- to 9-minute lectures on economic issues related to the Microsoft antitrust case. They then generated two arguments supporting the position that Microsoft is a monopoly and two arguments supporting the position that Microsoft is not a monopoly, based on the lecture content. After viewing both lecture segments, they had a 10-minute review session and were given 5 minutes to complete the assignment by annotating the lectures with their four arguments.

In the solo condition, participants worked alone on the exercise. In the asynchronous and live group conditions participants worked in pairs. To simulate students who meet face to face, pairs in the live group condition viewed and annotated the lectures together at a single computer, talking freely while watching, making annotations, and generating their four arguments to be posted through the annotation system.

The asynchronous group condition simulated distance-learning students who are unable to meet. Pairs worked asynchronously, communicating comments on the lecture material to their partner through MRAS. Each participant watched and annotated one lecture video while their partner watched and annotated the other. They then switched lectures. As they viewed their second lecture they could see and respond to their partner’s comments made earlier. During the 10-minute review session, each participant could see their partner’s responses to their comments on the first video they had watched. In a real distance learning class, a group might use email to discuss their choices. We approximated this by giving them 5 minutes to communicate using an instant messenger to reach agreement on the four arguments to be submitted as their final assignment.

After the task, all participants completed a 10-question questionnaire and 11-item comprehension test on the lecture material. The survey addressed feelings of engagement and enjoyment, satisfaction with work product and process, and sense of group cohesion.

## 7.2 Results

These results are reported in greater detail in [19]. In all conditions participants completed the task with little difficulty, found the lecture and the exercise stimulating, and reported minimal boredom and attention drift. There were no significant differences among conditions on these measures or on the comprehension test. The ease with which groups completed the task suggests that the system can support group exercises when students cannot meet face-to-face. Annotations were focused on the tasks: Some addressed the lecture content; others addressed the assignment., for example, “This is a good point. We should use it as an argument that...”).

| # Annotations          | Asynch | Live | Solo |
|------------------------|--------|------|------|
| Viewed 1 <sup>st</sup> | 4.67   | 2    | 2.17 |
| Viewed 2 <sup>nd</sup> | 5.33   | 2    | 1.67 |
| Review                 | 2.67   | 4    | 2.83 |
| Total                  | 12.67  | 8    | 6.67 |

Table 3: Average number of annotations per person.

As shown in Table 3, participants in asynchronous groups made significantly more annotations. They did so when viewing lecture material before their partners had made comments and also when viewing lecture material that had already been already viewed and annotated. In this latter condition about half of their annotations were replies to a partner’s previous comments, half were original comments. Consistent with our first study, people left more annotations when viewing material previously annotated than when viewing for the first time.

The group process, whether live or asynchronous, led to higher ratings of subjective work experience than did working alone. Live groups reported the most subjective satisfaction with overall work product, whereas asynchronous group members felt most strongly that they learned a lot from the exercise. This was consistent with our observations: The asynchronous group's very active instant messenger exchanges in the 5 minutes for discussing conclusions (averaging 18.3 entries) were task-focused; the live groups also discussed actively, verbally, but were more likely to engage in off-topic conversations that appeared to disrupt attentional focus. Live groups reported lower scores on whether the exercise made them think and analyze the lecture material, which may reflect this tendency to go off-topic. On the other hand, the asynchronous groups rated themselves lower on resolving differences and reaching agreement, a finding consistent with the literature on asynchronous communication tools.

The study also identified features to consider adding. Some groups wanted to move annotations directly from their personal or group annotation set to the class discussion set for their final report, which the system did not support. Also, some wanted to present their conclusions, two arguments on each side of the debate, as a single report rather than as distributed comments. This suggests the need for shifting easily to a "playlist" of annotations embedded in a general report, so readers would get a single block of text from which they could link directly to the four relevant segments of video. Students would also benefit from a detailed table of contents links to aid them in returning to specific points in a video in reviewing or preparing their reports.

### **7.3 Discussion**

This restricted laboratory study provided evidence that group exercises conducted with the asynchronous annotation tool could have the desired effect of engaging students with the material and each other, countering the procrastination and lack of engagement that have been major obstacles to on-demand education. The study also identified some capabilities that could enhance this experience. Of course, it remains to be tried in actual classroom contexts.

## **8. Conclusion and Future Directions**

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. Our approach has involved an iterative cycle between laboratory studies and deployment in actual courses. The system has performed reasonably well at each step, meeting the basic requirements according to the participants and from our observations and comparisons with control conditions.

The laboratory studies have led to changes in the design of the system and identified patterns to look for in actual deployment. In the field study, 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, instructors felt less fully engaged. Although we can look to ways to extend the system to provide benefits to instructors, a new technology is likely to the development of new processes of use, and adaptation will not be without challenges. In these studies we were using multimedia closely modeled on current live lectures; in the future, videotaped lectures may not be the most effective medium for asynchronous education, and this could lead to further changes in the technology and process of using it.

This research project is continuing. We are conducting ongoing field studies examining the use of text and multimedia annotation in university courses; comparing annotations linked to context with general online discussion boards; looking at the use of notifications to enhance annotations as a means of asynchronous collaboration; and considering different ways of using embedded, pre-authored self-assessment quiz questions in annotation sets as another tool to improve the educational experience for students. Further results will be posted at <http://research.microsoft.com/coet>.

## 9. Acknowledgments

Scott Cottrille, Brian Christian, Nosa Omoigui, and Mike Morton contributed valuable suggestions for the architecture and implementation of MRAS. The Microsoft Usability Labs provided the use of their lab facilities. David Aster, Barry Preppernau, Catherine Davis, Carmen Sarro, Shanna Brown, David Chinn, and Steven Lewis helped in conducting the C Programming Language courses. JJ Cadiz, Duncan Davenport, and Gavin Jancke provided technical assistance to the laboratory studies. Thanks to AJ Brush, Ross Cutler, Mary Czerwinski, Suze Woolf and Jonathan Cluts for suggestions.

## 10. 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] Barger, D., Gupta, A., Grudin, J., and Sanocki, E. *Annotations for streaming video on the Web: system design and usage studies*. Proceedings of the Eighth International World Wide Web Conference (Toronto, May 1999), 61-75.
- [4] Barger, D., Gupta, A., Grudin, J., Sanocki, E., and Li, F., 2001. *Asynchronous collaboration around multimedia and its application to on-demand training*. Proc. HICSS-34 Conference, CD-ROM, 10 pages.
- [5] Bates, A.W. *The Future of Education*, Teletronikk, Vol 92, No 3-4, 1996. pp 82-91.
- [6] 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.
- [7] Chalfonte, B.L., Fish, R.S., and Kraut, R.E. *Expressive Richness: A Comparison of Speech and Text as Media for Revision*, Proceedings of CHI'91 (1991), ACM Press, 21-26.
- [8] Computer Supported Intentional Learning Environments (CSILE) project. <http://csile.oise.on.ca/>.
- [9] 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.
- [10] CoWeb Projects, Georgia Tech College of Computing, <http://coweb.cc.gatech.edu/csl/29>.
- [11] CritLink Mediator by Ka-Ping Yee. <http://crit.org/critlink.html>
- [12] Davis, and Huttonlocker. CoNote System Overview. (1995) Available at <http://www.cs.cornell.edu/home/dph/annotation/annotations.html>.
- [13] Hypernix Goopy web communication system. <http://www.hypernix.com/>.
- [14] 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.
- [15] 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>.
- [16] Lawton, D.T., and Smith, I.E., *The Knowledge Weasel Hypermedia Annotation System*, Proceedings of HyperText '93 (Nov 1993), ACM Press, 106-117.
- [17] 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.
- [18] 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.
- [19] LeeTiernan, S. and Grudin, J., 2001. *Fostering Engagement in Asynchronous Learning Through Collaborative Multimedia Annotation*. S. LeeTiernan and J. Grudin, 2001. Proc. INTERACT 2001, ACM Press, 472-479.
- [20] Lotus Learning Spaces. <http://www.lotus.com/home.nsf/welcome/learnspace>.
- [21] 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.
- [22] Marshall, C.C. *Toward an ecology of hypertext annotation*, Proceedings of HyperText '98 (Pittsburgh, PA, USA, June 1998), ACM Press, 40-48.
- [23] 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.

- [24] 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.
- [25] Neuwirth, C.M., Chandhok, R., Charney, D., Wojahn, P., and Kim, L. *Distributed Collaborative Writing: A Comparison of Spoken and Written Modalities for Reviewing and Revising Documents*. Proceedings of CHI '94 (Boston, MA, April 1994), ACM Press, 51-57.
- [26] NovaWiz Odigo web communication system. <http://www.odigo.com/>.
- [27] Organik FAQ builder. [http://orbital.co.uk/organik\\_info/index.htm](http://orbital.co.uk/organik_info/index.htm).
- [28] Parker, A. *Distance Education Attrition*, International Journal of Educational Telecommunications, Vol 1, No 4, 1995. pp. 389-406.
- [29] 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>.
- [30] Scardamalia, M., & Bereiter, C. *Student communities for the advancement of knowledge*. Communications of the ACM, Vol. 39, No. 4, 1996. pp. 36-7.
- [31] 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](http://www5conf.inria.fr/fich_html/papers/P15/Overview.html).
- [32] 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.
- [33] Stanford Online distance education service. <http://stanford-online.stanford.edu/>.
- [34] ThirdVoice web annotations. <http://www.thirdvoice.com/>.
- [35] uTok web annotation system. <http://www.utok.com/>.
- [36] 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.
- [37] Vicarious Learning Project. <http://www.hcrc.ed.ac.uk/gal/vicar/>.
- [38] WebCT e-Learning Hub. <http://www.webct.com/>.
- [39] 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.
- [40] Wetzel, C.D., Radtke, P.H., & Stern, H.W. *Instructional Effectiveness of Video Media*. Hillsdale, N.J.: Erlbaum, 1994.
- [41] Zadu web customer contact system. <http://www.zadu.com/>.

## 11. Appendix

Figure 7 and Figure 8 show two variations of the MRAS interface. Figure 7 is the original interface used in the first laboratory experiments described in the paper. The displays are not all integrated into a single web page in this version. The URL at the top of the browser window identifies the video, the MRAS toolbar at the base of the browser window specifies which annotation server to connect to and which annotation set (questions, personal notes, table-of-contents, and so forth) to retrieve, and has buttons for performing “global” operations such as adding a new annotation. Once annotations are retrieved from the server, their headers (e.g., author and subject fields) are displayed in the “View Annotations” window, where they can be edited or deleted, and replied to. Beneath that window slides appear in the browser, although none were used in the study. The original interface allows greater flexibility in specifying annotation sets, but requires more knowledge and work by a viewer.

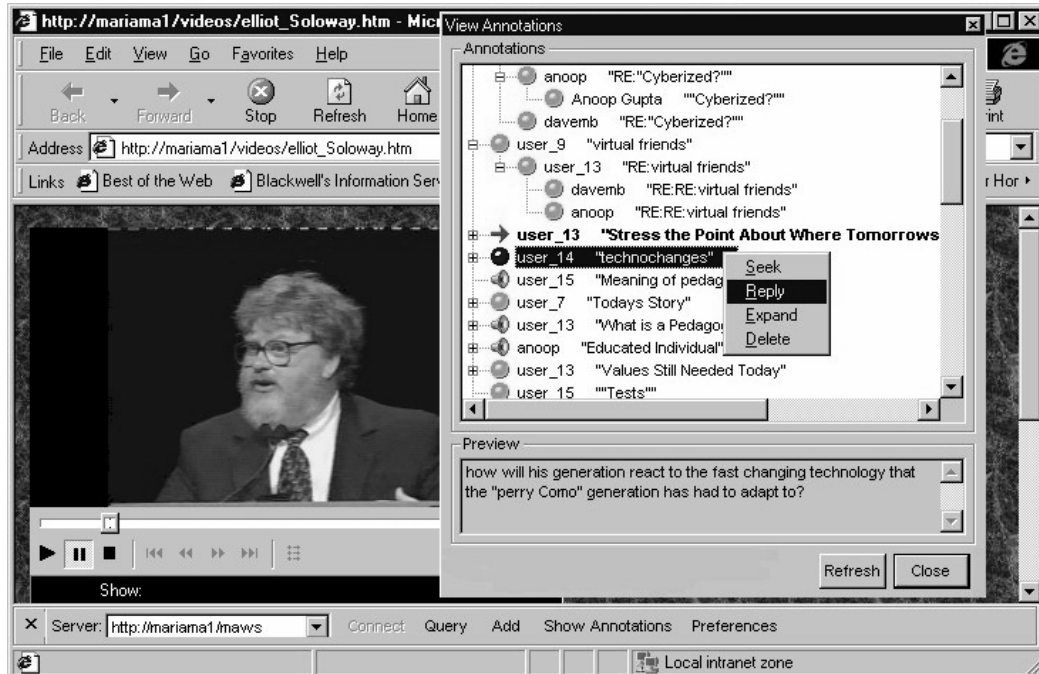


Figure 7: Original MRAS interface.

Figure 8 shows a more polished interface emerging from this work. It is clear that different interfaces will be useful in different contexts, which has led us to focus on building a set of components that can be used to create different interfaces through scripting.

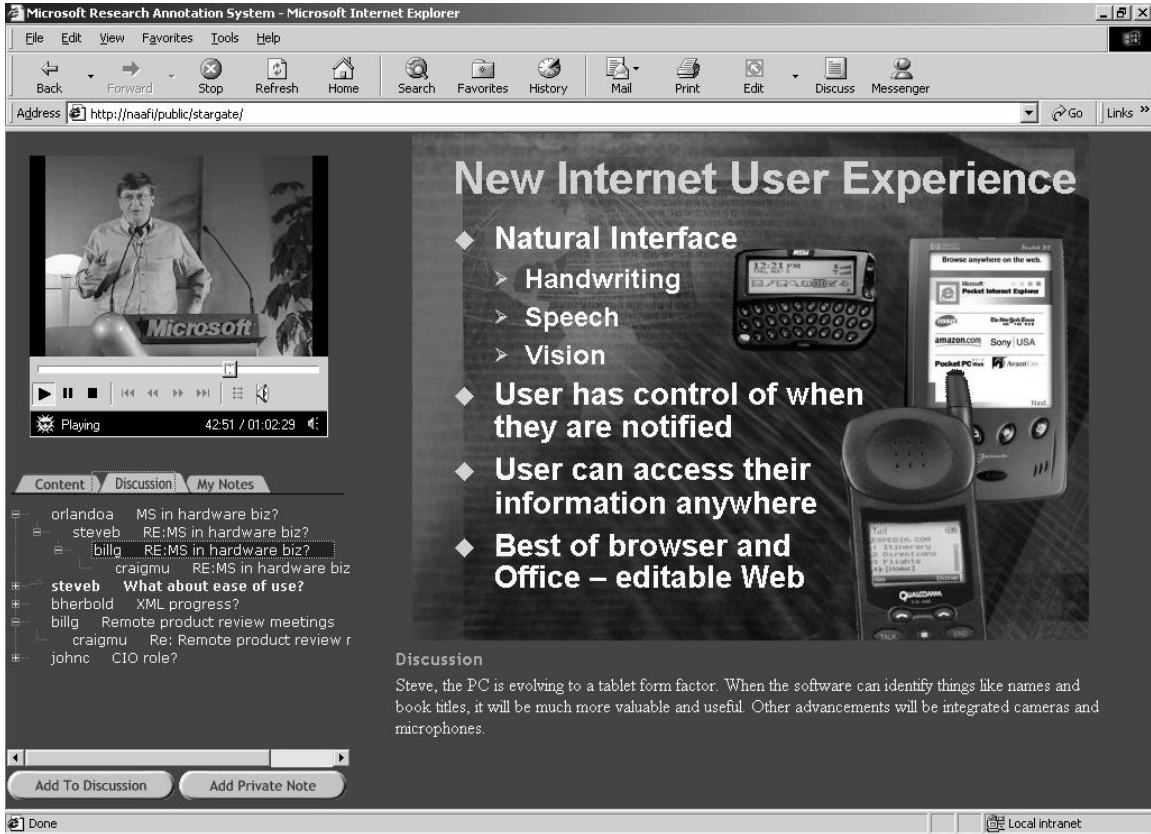


Figure 8: A more polished version of the MRAS workplace training user interface. This version uses the same web-based UI components as the version presented in Figure 2.