RIPE:
Rapid Instruction Production Environment

by Bruce Alan McHenry

B.S.	M.S.
Computer Science	Management of Technology
Massachusetts Institute of	Massachusetts Institute of
Technology, 1983	Technology, 1994

Submitted to the Program in Media Arts and Sciences
School of Architecture and Planning
in partial fulfillment of the requirements for the degree of
Master of Science in Media Arts and Sciences
at the Massachusetts Institute of Technology

September 1997

Author
Program in Media Arts and Sciences
June 13, 1997

Certified by
Nishikant Sonwalkar
Director of Hypermedia Teaching Facility
Thesis Supervisor

Accepted by
Stephen A. Benton
Chair
Departmental Committee on Graduate Studies
Program in Media Arts and Sciences

RIPE:
Rapid Instruction Production Environment

Bruce Alan McHenry

Abstract

This thesis makes two main contributions to the areas of distance learning and instructional content development. The first is the idea that multimedia interactive content should be created interactively by content experts in response to students' questions. To this end, a "mini-studio" for instantly publishing web pages from sketches and voice overlays was built. The second idea is that instructional content is similar to many products that are manufactured for mass distribution only after many design iterations which carefully consider customer needs. Hence, a quality improvement methodology was implemented in which ratings are collected from every student on every answer. The thesis hypothesizes that the ratings will help to distinguish, and then build upon, the most useful content.

Thesis Advisor:
Nishikant Sonwalkar
Director of Hypermedia Teaching Facility

RIPE:
Rapid Instruction Production Environment

Bruce Alan McHenry

The following served as readers for this thesis:

Reader
_____

Walter Bender
Associate Director
Program in Media Arts and Sciences

Reader
_____

Steve Lerman
Director
Center for Educational Computing Initiatives

Table of Contents

Table of Figures

Acknowledgments

I would like to thank...

Nishikant Sonwalkar for providing a supportive environment and many hours of discussion which have greatly helped to inform this work. This thesis would not have happened without Nish.

Walter Bender for his encouragement on this project and also SpeechWeb. He often seemed to be nearby with a guiding word.

Steve Lerman who provided advice and helped with the earliest proposal; Ron MacNeil for directing me to Steve.

Andy Lippman for being the first, over a decade ago, to support this general idea and admitting me to the Architecture Machine Group; Chris Schmandt for repeating the offer.

Linda Peterson for going an extra thousand miles in my case and also Santa and Laura for the cheerful receptions.

Whitman Richards and Steve Benton for facilitating my work with Nish.

Nicholas Negroponte for creating the Media Lab, a vision that has taken flight.

the students at the Media Lab who have inspired, encouraged or cajoled me including, but not limited to: Tomoko Koda, Jeff Herman, Sumit Basu, Rob Poor, Josh Bers, Marina Umaschi and Deb Roy.

Andrew Sempere for helping all around and being excited about doing this. I hope he fulfills his dreams.

1 Introduction

Colleague: Well, how did (the lecture) go?
Prof. Meyer: Terrible.
Colleague: What went wrong?
Prof. Meyer: I've forgotten what's hard.

[Minsky]

On-line teaching has been proposed as an alternative which will revolutionize education yet it fails to capture much of the human experience, the side conversations, the live opportunities for socialization and the many small gestures that bring a classroom to life. So far, distance learning projects have shown that distant learners perform about as well as those in the classroom. But, other than the geographic flexibility, there have been no outstanding advantages. [NJIT, HITE]

On the other hand, classroom based teaching does have intrinsic limitations. Teachers, even those who do remember "what's hard", cannot ensure that they've addressed the needs of every student in the classroom. Some students have seen the topic before, others are weak on a prerequisite; some have done all the readings, others aren't good readers; some have a special interest and want to delve into particular details, others only need an overview. Moreover, a typical class has many errors, omissions, disruptions, delays, and instances of poor diction and syntax.

Since we've had little experience with anything but the classroom approach, researchers are struggling to create something better with the suddenly available technology. This has led advocates to unwittingly recreate the 'invisible' limitations in on-line form. After refrigeration became available, suppliers of ice blocks used it to make ice used to supply home "ice boxes". The same phenomenon, where the initial uses of a breakthrough technology are initially used to augment the old approach, has been noted after the introduction of numerous new technologies [Utterback]. The real gains

typically require a revision of the infrastructure or processes or both. But first we need a new set of working assumptions. Table 1 provides some suggestions.

Old Assumptions	New Assumptions
Out with fixed times for regular lectures	In with lectures on-demand
Out with writing lesson plans and textbooks in isolation	In with using intensive feedback from students to develop content
Out with recitations and limited office hours	In with on-line tutoring all day (at least in the bigger classes) long
Out with separate times and places for lectures, recitations, office hours and home work	In with virtual locations where the different instruction modes intertwine
Out with the routines and the looking-glass of a few local faculty	In with star performers and the best ideas from faculty, TAs, students and especially alumni
Out with teaching classes	In with mentoring individual students
Out with fixed syllabi	In with smaller course modules that are more adaptable to interdisciplinary study
Out with routine classes	In with special events

Table 1: Towards a new set of assumptions

It does not work to change just one, or even a few, of these assumptions at one time. The changes reinforce eachother so that, while there may be great benefits from switching everything, piecemeal change will often make things worse. This phenomenon has been noted in many other instances of technology or process upheaval.

The objective of this thesis is to explore a methodology that is consonant with the new assumptions. The methods should ensure, above all, that customers' (i.e. students') needs are being met. This key idea leads to a focus on continuous, customer driven quality improvement. This idea was not adopted by American corporations until foreign competition forced the issue. While no such pressure yet exists at American universities, it will surely be noticed if and when the quality and cost of distance learning courses surpass those of classroom offerings in the best schools.

The methodology proposed in this thesis was conceived particularly to have the following characteristic:

Content must improve with use.

If better content attracts more students, the early adopting schools may reap advances that later entrants find difficult to replicate. Historically, revolutions in process and technology have displaced the former leading practitioners who were unable or unwilling to cannibalize their existing products or write-off their investments. [Utterback] We may be on the cusp of the most significant change in educational methods since Socrates. Top schools could be surpassed by worldwide teaching networks featuring star teachers and a commitment to 24 hour service. Those who do not lead, will surely follow.

Motivation

This work was originally motivated by the observation that undergraduates, including the author himself, were extremely reluctant to ask questions in large lectures even when the lecturer had "lost" a majority of the students. One response to this particular problem is to scatter audience response devices throughout a lecture hall. Portable response measuring equipment, which allowed students to select a number to show their level of comprehension, was used by MIT Prof. Thomas Sheridan and also at Cornell in the 1970s. The students and the teacher could see the class' continuously averaged comprehension index. This helped lecturers to notice problems as they developed and reassured students, when they were not alone in their bewilderment, to ask questions.

But there is a more fundamental problem. Especially in the oversubscribed lectures and recitations, there simply isn't enough class time to ask all the questions. Even if there was time, students wouldn't want to sit through all the questions. A good individual tutor, available to answer questions freely as the lecture progressed, would be much more efficient from the students' point of view

Some educational videotapes[1] suggest a way to supplant lectures and recitiations with something like on-line tutoring. The tapes can be more carefully structured and better finished than lectures although the production requirements are not complicated: a studio with one camera focussed on a pad of paper and, perhaps, another camera on the instructor

---

[1] The original idea for teaching using a pen, pad, camera and microphone came from videotapes made by MIT Prof. Al Drake (now retired) around 1980. The author found that, although Drake's lectures were excellent, the tapes, because they allowed pause and reverse, could be more helpful.

is all that is needed. The tapes have one key advantage over the lectures: they could be stopped and anything that was not understood can be repeated.

The simple studio, combined with easily manipulated digital media, suggests the possibility that the personal "tutors" could stand-by in their offices for questions from students. The questions could lead to exchanges which are like threaded discussions. The problem with threaded discussions is that they tend to include many poorly conceived or irrelevent which detract from the overall quality of the discussion.

Overview

The instructor has a workstation that is equipped to quickly create web pages that include sketches, verbal remarks and typed notes. Students view the instructor's presentation and ask questions. The instructor picks a student's question, usually from the FIFO queue, replies and publishes the answer for all students to see. E-mail notifies the questioning student that the answer is ready.

Students are asked to rate how well the question was answered. The ratings are averaged and used to rank the list of available questions.

Figure 1 shows the physical set-up.
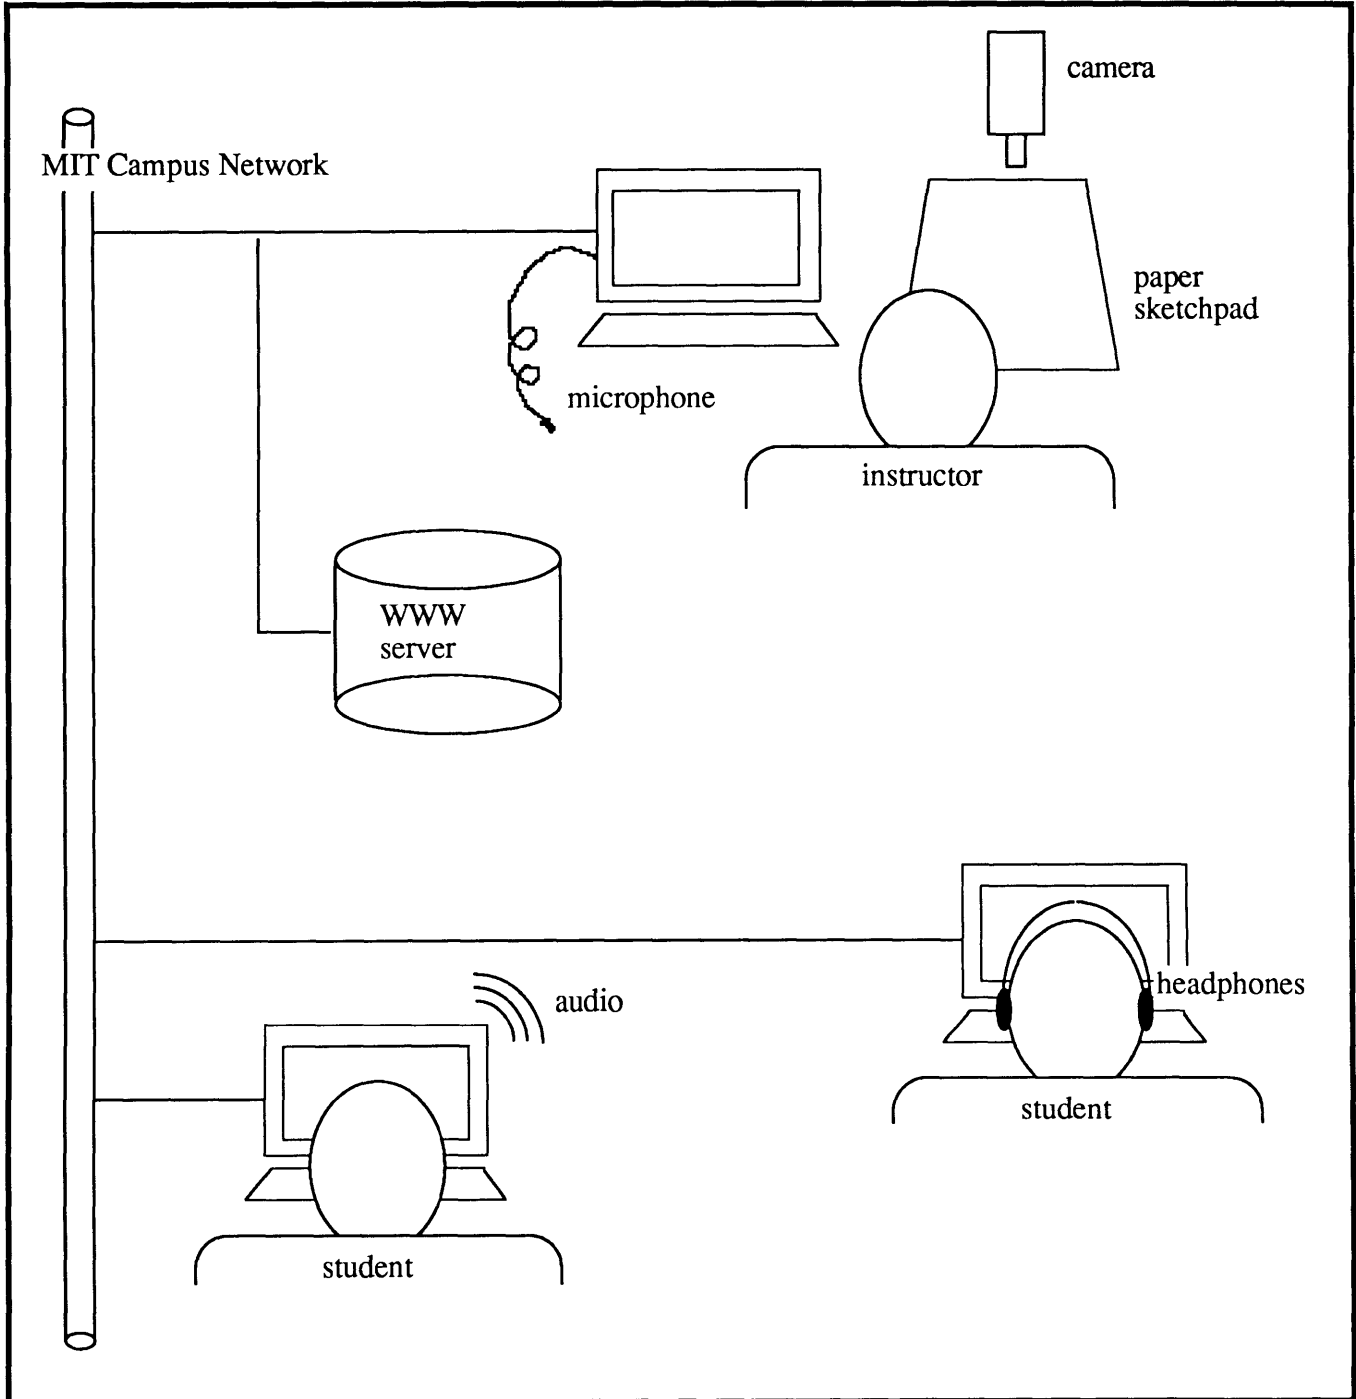
Figure 1: Physical Set-up

The methodology is the most important aspect of RIPE. The philosophy of Total Quality Management (Deming 94, Shiba 93) is a greater inspiration than pedagogy (which should not be ignored but treated warily to the extent that it has based on the old assumptions). TQM encourages us to acknowledge that the student is the customer from whom we should seek detailed feedback on the usefulness of the *educational product*. Figure 2 is a schema of a development process which is designed for quality improvement.

Each new cycle refines the content. Cycle often.

Students' questions guide instructional approach. Students provide explicit, fine-grain feedback.

design

test   prototype

"Sketch and talk" presentations can be prototyped quickly by using camera and microphone equipped machines.
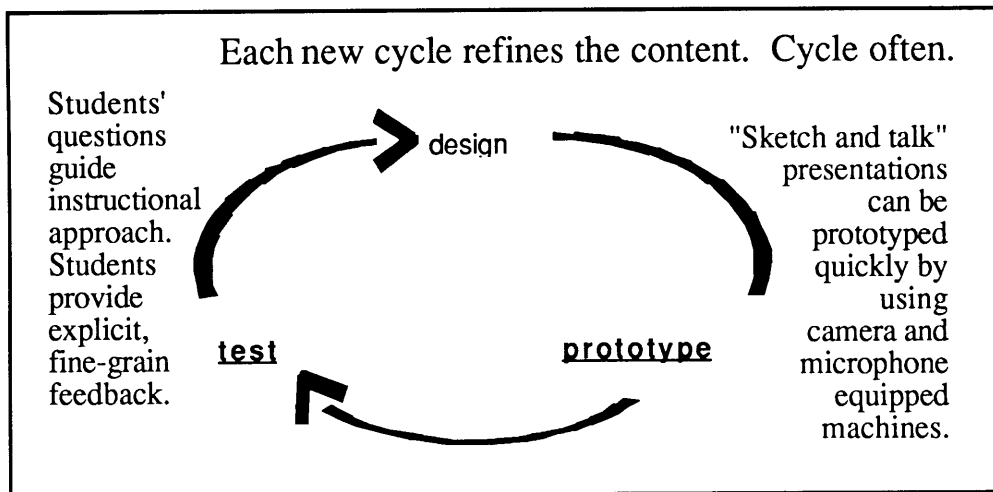
Figure 2: RIPE Content Creation Methodology

RIPE's focus is on prototyping and testing. Although individual prototypes may be rejected in testing, frequently repeating the cycle generally results in higher quality. RIPE is designed to reduce the cycle time and improve the odds of creating, identifying and giving broad exposure to the most useful approaches.

Therefore, prototyping should be quick and easy. RIPE facilitates prototyping by providing the author with a "mini-studio" equipped with a camera, which quickly takes snapshots of drawings and pictures, and a microphone. In order to make testing integral and not encumber the student, the same button both provides a rating and also navigates "back" from the answer being rated.

If students are customers, then instructional content, like any consumer product, must find needs to fill. Thus instructional development methodology may benefit from TQM's focus on the "voice of the customer". RIPE gives students two ways in which to express themselves: 1) as question raisers they identify their needs, 2) as reviewers, they indicate whether their needs have been fulfilled.

Several hypotheses are advanced to test the idea that the most useful questions and answers will be reused and improved, leading to an overall increase in the quality of the content. RIPE was briefly tested with the assistance of MIT Professor Feng using sample final exam solutions in couse 2.001.

Overview of the Document

The next section, related work, is divided into three parts representing academia, industry and quality management. Various implementation issues are discussed in chapter 3 followed by the users' views of RIPE's functionality. A short trial is mentioned in chapter 5, which also introduces several hypotheses which should be evaluated during the next trial. Chapter 6 seeks to identify market segments that RIPE is best suited to occupy. Then, future work is discussed both in terms of the long range and also the improvements that RIPE should include before the next larger scale test is undertaken.

2 Related Work

Academic Experiments[2]

Answer Garden

Mark Ackerman explores answers to the organizational memory problem [Ackerman 96]. Answer Garden 2 has two major functions. It helps to direct the users' queries using a process of sucessive escalation. The other major feature of AG2 is its collaborative refining. Both have significant implications for RIPE in its current and possible future incarnations.

The process of sucessive escalation means that AG2 might first direct the user to a FAQ database. If this does not answer the question, the system might then try increasingly compreshensive user groups. It will start with the user group that is deemed most likely to have experts that both know the domain and also the querying user. Finally, if the user remains dissatisfied with the replies received, AG2 might forward the question to a domain expert. The exact escalation process is configurable.

RIPE's escalation mechanism is simpler, and is entirely under the student's control. If the student does not find the answer in the immediate presentation, he can look at the list of questions (FAQ database) for that page. If this does not give satisfaction, he can escalate the process by asking the instructor a question. AG2 suggests an intermediate form of

---

[2] During the litterature search, only one reference was found to the idea of rating instructors' input. The quote, "...method of recording both accesses to the relevant pages and ways of soliciting responses from users on the usefulness of the pages." was found in a white paper from the City University in Hong Kong [CityU]. However the philosophy differs in that the following sentence, "Good comments will be included on the site while all comments will be forwarded to authors" assumes the existence of a moderator/censor as well.

escalation, such as posing the question to students who are on-line in the same class, that could be tried as a way to get a more prompt response. The use of intermediate escalation options helps to reduce the load on the domain expert or instructor.

Collaborative refining is defined as being comprised of four general activities: collecting, culling, organizing and distilling.

Collecting is a specialized activity in Answer Garden 2 becuase there are numerous information streams: NetNews, synchronous chat channels, distribution lists, and email. RIPE's collections process is much more uncomplicated because there is really only one information stream, the one in which students pose questions, instrustors answer and students rate. However, AG2 raises questions, "What if RIPE had other information streams? What would they be and how would they be handled?"

"Culling is a selection mechanism for identifying themes or threads that occur within a collection." In AG2, culling can be a collaborative exercise which "can further reduce the separation between users and experts, allowing for additional organizational learning." In culling, the less important or relevent material is ignored. With RIPE, this is not only collaborative but also continuous. Students are encouraged to rate everything (except the initial corpus) and the display of alternatives is adjusted with every new rating.

"Organizing allows one to group materials according to some classification schemes so to enhance their retrievability and understanding." Outlining, user-defined indexing, and keyword indexing are used to organized items in AG2.

AG2 has 2 distilling processes: 1) Manual concatenation of selected items, 2) Forgetting of older (presumably unaccessed) items. In neither case is material actually deleted, presumably because it's text based and there is no pressing need.

Ackerman does not want to make a clear distinction between experts and other users. He points out that regular users may assemble FAQs as a way of learning, or that a particular distillation may be designed to match the expertise of the group. The expert distillation is not necessarily the most useful one. [Ackerman email]

Paper-Based Audio Notebook

Lisa Stifelman has implemented a physical world (notepaper and pen) instead of a WWW instantiation of the audio/visual interface [Stifelman 96]. Designed specifically for note taking, this audio notebook has a speaker and microphone. It accepts a 5.5-by-9 inch paper pad which has marks on it to code the page numbers. While recording, the device correlates the page number and the users' writing with the audio so that, on playback, the user can touch annotations in order to trigger the corresponding playback.

The main novelty compared with [Whittaker '94] (see below) is the choice of paper as the input medium. Stifleman also notes that paper can easily be "torn out, stuffed in a pocket, handed to a friend and generally put in a place where it can be found when needed". Although Stifelman does not say so in these proceedings, tablets made with digitizing overlays on LCD screens suffer from poor resolution. Similar inadequacies with digitizing tablets led me to use paper as the input medium for RIPE.

During the usability studies, Stifelman tested the device with a subject who tried to start and stop the recording to capture only the interesting portions. This turns out to be futile because you usually don't know that something interesting is being said until afterwards. However, it led to the interesting observation that others noticed when the recording was being halted. This was particularly true of the speaker who said, "why are you pausing me!" The implication for RIPE is that people are interested in whether they are being listened to or not. We might infer that teachers will be interested in knowing what students are interested in hearing.

The audio notebook provides two ways to specify the time of playback. One is to touch the page where a note was taken, causing audio to start at that point. The other is to use a slider along the side of the page. At first, users thought the slider was redundant. However, they later found it to be a useful way to "run down a page". Certainly, it would be useful for locating speech that occured during periods when no notes were taken.

One could imagine future versions of RIPE with large images, perhaps some simple animation, and audio lasting many minutes. Both of Stifelman's techniques could be used to help solve the problem quickly finding the relevent section of the audio.

Filochat

Steve Whittaker et al. developed a device that, as the predecessor to Stifelman's, also integrates handwritten notes and recorded audio [Whittaker 94]. Tested with 67 users, the study showed objective benefits over notes alone. Like Stifelman's device, this one featured audio retrieval based upon pen position. This access method was shown to have better performance than current products such as dictaphones. Users perceived that they were able to prepare better minutes of a meeting by using Filochat. An unexpected consequence of the trial was the discovery that this device could be used as an audio editing tool.

The researchers found that users would use idiosyncratic symbols in order to note a section of audio. They thought about developing functionality around that used this: "One possibility is that the system could provide a limited range of user definable symbols to enable the identification of key "sound bites" without the need to take detailed notes. Thus a tick, question mark and cross could be used to access quotes that the note-taker agreed with, did not understand, or disagreed with."

Whittaker has also suggested the use of intonational analysis to determine affect or identify segments as well as keyword spotting [Whittaker 93].

Simultaneous On-line and Textbook Development

Julie Rattner is concerned with "a critical issue (that is) the simultaneous development of paper and on-line versions of the (engineering) textbook... Many of these problems were resolved swiftly with human factors input using templates, style guides, and iterative usability testing of both paper and on-line versions." [Rattner 96] RIPE shares Rattner's interest, especially in iterative usability testing. Rattner also raises the possibility of using intensive on-line testing in order to create better textbooks.

Importance of Drawing Tools

[Farinetti et al] This distance learning experiment was conducted at 128 Kbit/sec with videos of a chalkboard, audio and application sharing. The authors reported that students noticed minimal differences between remote and on-site instruction as long as there were prepared materials. From RIPE's perspective, the most important finding is that better input tools are needed for hand drawing in response to unanticipated questions (video images of chalkboard were too poorly defined). Since good visuals could not

be created on-the-fly, questions had to be anticipated and this entailed more work for the tutors.

Audio Clip Length

In the absence of better techniques to access the audio, there must be an optimum audio length. If the soundtrack is too long, it is impossible to find anything; too short and it's impossible to say anything substantive. [Gloor et al] found that sound clips for web based publishing should be about 30-50 seconds per page with multiple clips per page. (Although the work of Stifelman suggests that audio could be much longer with a better paradigm for random access.)

Real World Learning Systems

Teleconferencing Marketplace

The commercial distance learning expenditures are currently dominated by videoconferecing technologies. When positioned as an alternative to jet travel, it is possible to justify very large expenditures for high quality images over leased T-1 and reserved satellite channels. Often the systems will find mixed use, both for meetings and corporate training.

However, the value of synchronous "talking head" capability applies mostly to personal introductions. For actual work, asynchronous communication focussed on the task, not the faces, is preferable [refs Xerox PARC], Although this has been known for over a decade in the Computer Supported Cooperative Work (CSCW) community, teleconferencing continues to dominate (mostly corporate) distance learning expenditures.

There are signs, however, of a changing tide. At the most recent International Distance Learning Conference, a featured speaker at the general session, Sir John Daniel, Vice-Chancellor of the Open University, blasted the participants' focus on "students in remote classrooms looking at TV screens". Drawing a distinction between individual learning and group teaching, Daniels criticized the latter saying that the real challenge is to reach many more people: inexpensively and conveniently. (The Open University has begun to experiment with a few web based courses but most of the 150,000 enrolled students use TV, tape, books and local study groups. [Open University])

There are at least three companies selling virtual classrooms on the web [Avalon, Centra, iFactory]. All three have substantially modelled their products on the remote video model and positioned themselves as a more cost effective alternative. They provide synchronous audio and images, along with floor control. In addition, in order to mitigate the loss of a face-to-face contact, they have implemented several forms of feedback.

Two kinds of feedback (pace and comprehension / confusion) are for use during presentations. Students may select one of a number of options (i.e. too fast, OK, too slow; I'm confused, OK, I'm bored). Students provide feedback only when they think to do so. Feedback is not, as it is in RIPE, part of navigation. Hence students tend to forget about the settings until they want to register a problem and then forget to readjust back to normal later. (The programs tend to let student feedback renormalize after a fixed time.)

These system also has a mechanism by which a student can "raise his hand". The instructor can yield to the student, effectively giving him control over everyone's audio and display.

Lotus Domino

Domino is an applications and messaging server with an integrated set of services that help to create secure, interactive solutions for the Internet. This is a alternative platform, as compared with a Netscape server, that could be used to develop web based applications. The relevent features include object storage, and integrated messaging. It could be used to construct a course builder, or a future version of RIPE.

Course Builder Applications

Many schools have developed their asynchronous on-line course in the context of frameworks which provide administrative functions and some facilitation of the authoring process. Some of the more significant ones are cited below.

HITE

Developed at the MIT Hypermedia Teaching Facility, which was started in 1993 and is under the direction of Dr. Nishikant Sonwalkar, the Hypermedia Instruction and Teaching Environment [HITE] supplemented five classroom based courses in mechanical engineering during the 1996-1997 academic year. The content is WWW readable and was built using textbooks, re-

purposed videos and virtual laboratories. With the exception of a few video related problem sets, the HITE materials are supplemental to courses that are being taught on campus. According to the most recent [HITE] report, students were most appreciative of the virtual laboratories and the videos. Over the course the term, the percentage of students using the other on-line content (derived from a textbook with the addition of hints to problems sets) declined while the number of acceses remained relatively constant. This suggests that a small core of students (about 5%) would use on-line text despite ready availability of books and classes.

WebCT

"[WebCT] was developed in the Department of Computer Science at the University of British Columbia. The faculty member in charge of the project, Murray W. Goldberg, has had experience building, delivering and studying the success of web-based courses, and of web-based material used to supplement existing courses... WebCT also provides a wide variety of tools and features that can be added to a course. Examples of tools include a conferencing system, on-line chat, student progress tracking, group project organization, student self-evaluation, grade maintenance and distribution, access control, navigation tools, timed quizzes, electronic mail, automatic index generation, course content searches and much more."

CyberProf

"The University of Illinois' CyberProf is an 'intelligent' (quotes added) student-computer interface and consists of several modules: grading system, gradebook, network TA problem set editor, conferencing, survey manipulating utilities, lecture note editor (The HTML lecture note editor has been adapted to accept modified LaTeX tags.), index maker, plotting and drawing." Approximately 8 courses at UIUC use CyberProf authored courses and there are several more courses offered at a community college and a high school. [Cyberprof]

FirstClass

This is a commercial, multiplatform email and group collaboration product. The Open University, Europe's well-known distance education institution, has been using SoftArc's FirstClass to administer courses to remote users for several years. With 15,000 users in 1996, the site is now planning to double its student base to 30,000 users in 1997. FirstClass provides email, realtime chatting and threaded discussions; it is positioned as a scaled back version of Lotus Notes which is easier to install, requires less computing power and less training than Notes. [FirstClass]

Authorware

There are also authoring systems for content creation. Authorware, is designed to "allow subject matter experts, not programmers, to create compelling interactive courseware." The interface is based on iconic flowcharts which allow the user to layout a interactive solutions for web or CD delivery [Macromedia]. This is often used for developing drill and test materials.

PageMill

This is not specifically a product for education content developers but, like RIPE, it facilitates the creation of web pages. PageMill is probably the best of several competitive products [Adobe]. It offers many more features than RIPE but, unlike RIPE, does not directly control a camera and microphone. Hence, page publishing is slower and more complicated than in RIPE. The applications are incompatible because RIPE generates pages on-the-fly.

Hypermail

"Hypermail is a program that takes a file of mail messages in UNIX mailbox format and generates a set of cross-referenced HTML documents. Each file that is created represents a separate message in the mail archive and contains links to other articles, so that the entire archive can be browsed in a number of ways by following links." In essence, Hypermail creates links between each mail message and the message it replies to as well as next message in the thread.

Webcrossing

This is a commercially available application that adds discussions and chat groups (unthreaded), using text only, to web sites. [Webcrossing]

Total Quality Management

RIPE's approach applies some principles of Total Quality Management (TQM), which is widely practiced in other kinds of product development, to the development of educational materials. TQM was originated by W. Edwards Deming, a statistician who was recruited by Gen. MacArthur to help the Japanese rebuild their economy. One of his observations was the importance of shortening the prototype and test phases of the development cycle so that there could be more cycles. Since the product improves with

every cycle, increasing the number of cycles will increase the quality of the product.

While the origins of this thesis came from frustration with inaccesibility of instructors and the "impedance mismatch" between professors and undergraduates, the TQM philosophy provides another source of ideas and inspiration for RIPE [Shiba 93][Deming 94]. Improvement is seen as a problem solving process. Hence, the effective way to improve quality is to improve the process used to build the product. The focus is not on the results, per se, but on the process used to guarantee good results.

Management by process, as applied to RIPE, would work as follows: First, set a goal (develop an on-line course) and develop an implementation plan (schedules, syllabus, lectures) Then, develop a system for measuring results (students' ratings, questions and test results) and undertake the task (creating multimedia educational content) while monitoring the results. Use the results to figure out why the plan didn't work and revise the plan, implementation and measurement goals. Repeat as required.

This is restated in the Plan-Do-Check-Act cycle:

Plan: *Pick the problem that is most responsible for poor results, analyze the root causes of the problem, and plan countermeasures to fix its root causes.*

Do: *Do the improvement*

Check: *Was the improvement effective? Why/why not? (Deming later renamed this step 'Study'.)*

Act: *Incorporate the improvement, throw it out, or try the cycle again.*

RIPE closely mirrors the PDCA cycle. Students directly participate in the 'Check' step and determine whether or not the improvement was effective. TQM is highly critical of what is called a "product out" mentality. This is the idea that a product is created in a kind of vacuum and then delivered to the marketplace without collecting input from customers about their needs.

3 Implementation

Software Architecture

The software consists of 12 PERL scripts. The largest of these modules (ShowPage.pl), is responsible for presenting a content page with its graphics, text and audio. ShowPage.pl displays different controls on the page depending on whether the user is a student or instructor.
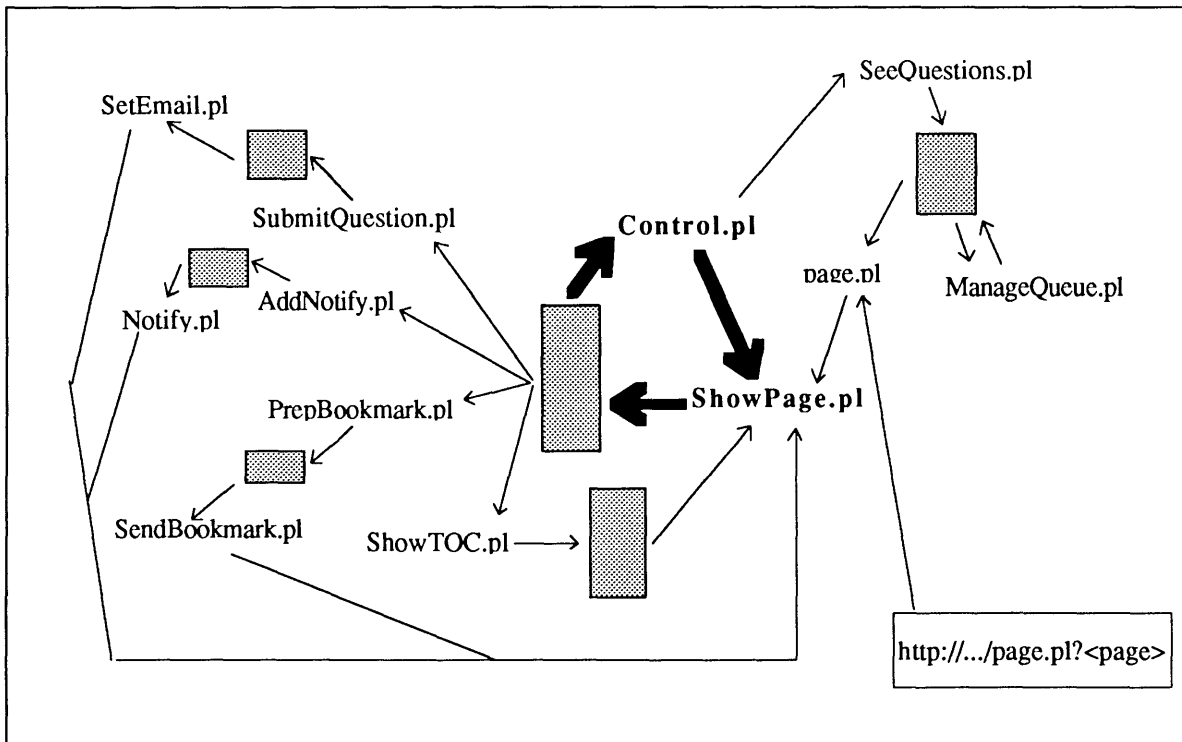
Figure 3: Relationship of PERL scripts to WWW pages

Most of the buttons on a content page call a control module (Control.pl) which implements the majority of authoring and student functions. (take/delete a picture, record audio, update the text, publish the page and notify the interested parties, jump to previous/next/contextual page, rate an answer).

The other scripts provide supporting functions, such as allowing the user to confirm their questions and email addresses, display the queued questions, send themselves bookmarks, add themselves to an answer's publication notification list, and show the table of contents.

Design decisions and tradeoffs

Single Author

This version of RIPE supports only one author since only one author was expected for the trial. This decision eliminated the need to tag answers by author and lock pages to prevent modification by more than one author at once. However, it may be desirable in a larger scale trial to post the same question to more than one author (some of whom might be students) and allow the answers to compete (see future work).

'Published' Answers Remain Modifiable

In the early implementations, pages could not be modified after being "published". This was done so that questions would not lose their context if the referred to material was removed. This was changed so that mistakes on published pages can be corrected. A problem remains: the newly contextless questions have to be identified and edited out. If they were well rated, they will probably stay high in the rating based rankings. (This favors the argument that the rankings should be based on recent level of use, rather than average score.)

Work in Progress is Visible

The instructor's work(s) in progress is visible to the students. They could instead be hidden until published but it was felt that there is little harm in allowing this access.

Manual Page Advance

The early design would have allowed the main presentation to advance automatically. i.e. The end of the audio track for each page would commence the download and playback of the next. However, this ideal of continuous playback gradually lost currency. Some pages will have no audio by which to set the display interval. But the main problem is that users need time to ponder the material and decide whether there are any important questions before advancing. Even if there are none, an obligatory interaction that requires minimal attentiveness helps to ensure that students, while they can still fall asleep during lecture, won't miss much if they do.

Voting also Serves as "Back" Button

In order to encourage users to rate the answers, the act of voting also returns users to the origin of the question. Unfortunately, this does not always work because users expect to use the browser's "Back" button in order to return to the origin of the last hyperlink, or question in this case. (See trial results for an estimate of the frequency of occurence.) There is no way to interdict the Netscape browsers' 'Back' button. The browser toolbar can be turned off but then some users are in the habit of going back by using a menu off the alternate mouse button, or the "Alt+Left" key combination. However, it is possible to intercept "Back" button in the CGI script and pointedly ask the user to either vote or explicitly abstain from voting.

Unrated Top-Level and Individual Pages

During the early design phase, it was decided that the top-level pages would not be rated, partly in order to accomodate the possibility of advancing the pages automatically. Now that manual page advance is considered preferable, one could revisit the decision to not rate the top-level pages or the individual pages that constitute answers.

Rating the top-level pages are not very useful since there are no alternate paths. Although ratings may give authors a quick indication of problem pages, the questions will suggest more explicit ways to improve the pages. While it does not presently seem worthwhile to collect ratings for every page, such extra ratings may be useful the future when detailed student profiles and alternate pages exist.

Selection of What to Rate

There are two situations where ratings of questions and answers could be helpful: 1) The student knows that he doesn't understand but is not sure why. 2) The student has a question in mind and there are several similar or identical questions. The first situation implies that ratings should be based upon replies to the question, "How useful was this answer?" The second situation needs an answer to, "Was the question answered?"

The first situation is certainly not uncommon, and at least one professor has said that this kind of non-specific question is the most typical[3], at least in undergraduate courses. The alternate view is that students more often have pre-formed or nascent questions that are likely to be similar to the questions other students have. RIPE asks students this latter question but the discussion is not closed (see Future Work).

Question Ordering within the List

The questions are ordered, or ranked, based on average ratings. The main alternative would be to rank questions based of the number of uses. Both approaches have merit. The average rating approach allows a new question to jump to the top of the rankings where it is more likely to be seen and used. Although not a feature of RIPE, a forced initial bias (e.g. automatically scoring as if several top ratings had been given) will help to ensure that new questions get an opportunity to compete with the other top ranked questions. Since the answers are not static, more accurate rankings could be obtained by biasing for recent ratings.

Ranking by use, on the other hand, would give prominence to the questions that the most students have used. This will confer advantages upon questions and answers that have been in use for the longest time. The older materials will have had more chances to accrete clarifying or interesting follow-up questions, even though the original answers may not be the most clear and direct. Furthermore, if material tends to improve with use as conjectured in the hypotheses section, ranking by use will create a positive reinforcement cycle which would make it very difficult for fresh new answers to displace the old standbys.

---

3 It presents the idea of a student interface that uses a thumbwheel: one direction would signify "tell me more, tell me more" and the other, "tell me less".

On-the-fly Web Page Creation

There are no stored ".html" pages of instruction; HTML output is created on-the-fly from datafiles. This choice was dictated by the fact that each page can have four views: 1) unpublished/author, 2) unpublished/student, 3) published/author, 4) published/student. The dynamic page creation entails some compute overhead but this is unremarkable compared with the time required even for local downloads of the graphics in the web pages.

Flat Files

A flat file database, with one file for each web page, consists of lines of tagged descriptive data. While this is fast enough for this application, future design changes may force adoption of a database arrangement with faster access. e.g. The flat file arrangement may lead to unsatisfactorily high search times, especially as more detailed records of student's voting patterns are kept. At present, only the last vote and voter ID are retained and checked. This provides a first level safeguard against double-counting that works in the most common instances when a user 'reloads' the browser immediately after voting, or jumps back and forth between a question and its context.[4]

Local E-mail Bindings

RIPE develops its own database mappings from user IDs to e-mail addreses. When there is no knowledge and the user performs a function that sends e-mail (New question, bookmark), he is asked to fill out his address. On subsequent invocations, the student's address is filled in from the database. The alternative would be to set up e-mail addresses when the students are registered as users.

CGI "Get" Convention

The Common Gateway Interface (CGI) "get" method of passing information between client and server was used. The main advantage, for debugging, is that the arguments are visible in the "location" window of the browser. One disadvantage is that the arguments can be changed by a mischievous user, causing unpredictable results and possible database corruption.

---

[4] Someone who was determined to skew the ratings could periodically recast their votes (as long as their vote was no the last one received, it would count again) However, the votes are logged so an active administrator will be able to readily apprehend the perpetrators. In any case, students have little incentive to rig votes on the instructor's content (but this will change if students become contributors and, perhaps, get real credit or cash for heavily distributed submissions).

Rating System

RIPE claims to make quality control integral to the content development process. This claim is based largely on the rating system. Yet several potential problems merit discussion.

Attractive but Simplistic Answers

The main objection is the too simplistic answer problem, in which students favorably rate an attractive solution that actually fails in some critical but non-obvious way. If students ever realize that the answer is bogus, it is long after they have seen and reviewed the answer. Under a system using student ratings alone, there is no redress: the misleading but highly rated answer will remain highly rated, pulling students in and providing them with a false sense of satisfaction.

RIPE avoids this problem because it allows only one content creator and that person is also expected to be the content expert. As an expert, they are expected not to provide misleading answers. However, reliance on a sole authority can be limiting when the authority is not readily available. In this case, material might be authored by teaching assistants and even students themselves. How is the material to be refereed then?

One could ask the content expert to review the newly created materials. It takes less time to review than to create. There could be several options in the event of an incorrect answer. The first, which is supported by the current version of RIPE, would simply be to edit the error (and provide notification to the author and the misled students). Another approach would be to leave the material unmodified but to have a specially weighted rating for the content expert - editor. This instructor rating could, for instance, carry as much weight as all of the student ratings.

However, the preceeding solutions will not solve the problem if the content expert doesn't have the time to review all of the material. Even if they do eventually get around to it, many students may have already ingested the error.

While RIPE assumes that there is a content expert whose authority is unchallenged, the experts are not always so readily available (This is especially true in the newer and more advanced courses where RIPE is most likely to be useful). How can a system reduce the load on the principal content expert, encourage teaching assistants and students to contribute,

and also quickly catch bogus answers? The question begs for a upgrade of the model and is discussed later (see Future Work).

Unfair Initial Ratings

The first ratings received, if low, will make an answer drop to the bottom of the list. Even though the initial ratings may not be representative, the answer may then languish unnoticed. One solution would be to showcase a new answer by pre-loading its rating with a (high) initial bias that is gradually repalced by real ratings.

Historical Burden

RIPE averages ratings over the entire lifetime of each answer. If an answer receives poor ratings at first but is then improved, either by modifying the answer directly or by new sub-questions, it will still carry the burden of its earlier scores. Since new material will affect ratings, we could simply forget the oldest ratings.

The converse situation, overly good ratings initial ratings that decline, will be self-correcting at a small cost to the students who are inappropriately exposed to the material.

No Well Defined Cutoff

One would like there to be a mathematically well defined cutoff point, below which answers could be deemed unfit for use and deleted from the list of answers. RIPE does not implement such a cutoff because it is impossible to guarantee that an answer won't be valuable to anyone in the future.

Students will use their own judgement to decide where the cutoff lies. However, as a practical matter of storage space limitation, it will become necessary to delete some materials. While the main criterion should be ratings, the amount of use may also be considered. A heavily trafficked question, though poorly rated - this should be rare - would be retained until better material is produced. Also, a well rated answer that only a few people have asked, may have to be sacraficed to create space for more mainstream content.

Integration with HITE / Logging

HITE is the Hypermedia Instructional Teaching Environment in which RIPE was embedded. RIPE is referenced from the HITE home page and, since it

shares a web server with the HITE pages, is able to share administrative groups for authenticating users.

At the time of the test, RIPE did not use HITE's statistics gathering facilities but logged its own events. The logged events included: page publication, page 'hits', and ratings.

The "look and feel" of the pages is not yet consistent with HITE.

Choice of platform

It is assumed that the network will be fast enough to download image files in reasonable time. This makes access by modem problematic since web pages, which typically will contain a standard image (60K bytes - from the camera), would download in about 16 seconds at 30K bps. This suggests that users should really have faster connections to the inter/intranet.

Audio, which streams at 20K bps, would work over a fast modem since the images and audio are downloaded sequentially. Together with snapshots from the instructor's camera, speech and stills create a "shared desktop" where ideas can be sketched out and discussed. Video is not supported primarily because of bandwidth/storage considerations and also because "talking head" shots add comparatively little value compared with the cost. Subject specific video clips could be linked in with some relatively minor modifications to RIPE's code.

The hardware environment was largely defined by the availability of equipment for server and client machines. A Silicon Graphics Indy workstation was chosen as an authoring station because of its standard video and audio capabilities. The Indy was also used as a web server which facilitated the software development: the authoring scripts can also directly invoke the image and audio system calls on the combined web server / authoring workstation. However, this is also significant because the authoring machine must also be running a web server. (Students could use a different web server as long as it shares files with the authoring machine via network file sharing.)

PERL is the generally preferred CGI scripting language and offers faster prototyping than compiled languages such as C or C++. The server was

never expected to handle a high volume of requests. The web server (initially Apache then Netscape) was chosen based on availability.

The client is controlled using HTML and JavaScript (which shares JAVA's name but not its key features).

Lotus Notes was also considered as a possible platform. Notes requires a proprietary client environment that is not available on Athena machines. Lotus' newest web server platform, Domino, solves the client side problem but was not available at the time implementation began.

Video stills are captured and converted to .gif files with stock tools supplied by SGI. The process takes rather a long time, up to 20 seconds and requires two steps. Unfortunately, SGI has no single capture and compress utility so instead the 1MB image has to be captured to disk and then compressed down to about 60KB.

A camera was used instead of a digizing tablet for several reasons: 1) Cameras are convenient. Video inputs are standard equipment on SGI workstations - even though the image quality of the stock camera is poor. 2) Graphics tablets have problems with aliasing and/or insufficient data rates leading to even worse resolution of handwriting than with the (poorest) stock SGI camera. 3) Tablets are not helpful for capturing supplemental pictures and graphs that have already been rendered on paper. A higher resolution video camera could be connected to the SGI workstation which would yield sharp images when opened in a typical width, 600 pixel wide, browser window.

The choice of audio format came down to a decision between RealAudio, a commercial supplier of streaming audio, and non-streaming solutions already available on the chosen hardware. Because of quality problems with RealAudio in December '96, we initially decided to use the audio formats indiginous to the Athena workstations (SGI's .aiff or Sun's .au format). However, by May '97, RealAudio had improved playback quality and reliability enough to be incorporated. RealAudio has several advantages: 1) it streams the audio data so that playback can commence almost immediately regardless of the size of the audio file, 2) it plays on almost any client with the RealAudio plug-in, 3) it checks for an existing audio control panel when opening new audio tracks so that it does not leave the user with a stack of control panels to close. It does not have one control which would be most useful: a button that would rewind approximately ten seconds. It is

commonplace to miss something that has just been said and a short rewind would really help.

4 User Interfaces

RIPE can be used concurrently by many students and one author.

Student Interface

The student interface (see figure 4) is designed to be simple and "intuitively obvious" to use. The page layout places the elements that are most likely to be seen and used (picture and text; page and rating controls) at the top of the page so that most pages can be viewed in their entirety without scrolling.

Navigation

Paging Controls

A presentation consists of a linear sequence of web pages. Two buttons, "Next Page" and "Prev Page" serve to go forwards and backwards. If on the first page, the "Prev Page" button disappears (as in figure); the "Next Page" button vanishes on the last page.
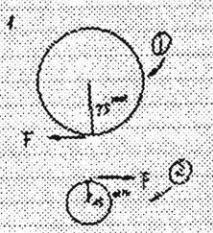
Pursuing Questions

The questions are ranked according to their average rating, highest first. The question text is a hypertext link to the page with the answer. Students can return from the answer page by rating the answer (see below).

New questions that haven't been answered yet are marked 'pending'. So long as the instructor has not begun to reply, the pending question's link points to a short form (see figure 5) which can be used to get notification when the answer is published.

Figure 4: Students' View

0 question(s) are ahead in the queue. In order to be notified when this question has been answered, enter your email address here:

nishikant@mit.edu

Submit   Abort

(Run "xbiff" at your Athena prompt so that you will be notified when your email arrives.)

Figure 5: Other students can also be notified about pending answers

Once a reply is underway, the student will see the unfinished work under a banner headline, "This page is really under construction." However, it is not possible to ask a question until the page has been published (see below).

Audio Controls

The audio control panel is activated the first time that audio is needed. The RealAudio panel has volume, play/pause, and stop controls. It also has a slider to randomly access any part of the audio segment.

Questions and Answers

Writing a New Question

Students may type a new question into the provided text box. After entering the question, the user is presented with a form (see figure 6) which allows them one more chance to read and edit their question. The form also solicits an optional e-mail address for the purpose of notifying the student when the answer is published.

The new question is immediately appended to the instructor's queue and also appended (marked 'pending') to the list of questions on the page where it arose.

Figure 6: Student Question Dialog Box

Receiving Notification of a Reply

A student will receive notification via e-mail when their question has been answered. (see figure 7) The URL that they are provided will take students to the context where the question is published. To reach the answer, students will paste the URL into their browser, re-orient in the question's context, find the question and link to the answer.

```
Date: Mon, 2 Jun 1997 21:51:59 -0400
From: RIPE
Subject: answer
Apparently-To: bru@media

"Why doesn't phi(t) = phi(1) + phi(2)?"

The answer is at:
http://nish.mit.edu/cgi-bin/RIPE/page.pl?pfa1.1:4.1
```

Figure 7: E-mail notification received when answer is ready.

Rating an Answer

When a student has finished with an answer, which does not necessarily mean that he has seen all of the pages, he is supposed to rate the answer on a scale of 'grades', A through F (Figure 8). This grade will be averaged with those of other students. (If user IDs are in use, a student won't be able to "stuff the ballot box" by voting more than once. Only a user's last rating is counted.)

The rating function serves the dual purpose indicated by the label on the rating buttons, "Rate overall answer(s) to the question above and go BACK". Where 'BACK' assumes the meaning which it normally has in web surfing: to go back to the context which contained a hyperlink to this answer. In the case of RIPE, every such hyperlink is a question.

Q: why is it 3 times phi?

Because the small wheel will rotate three times when the big wheel rotates once.

Rate overall answer(s) to the question above and go BACK [A] [B] [C] [D] [F] [abstain]
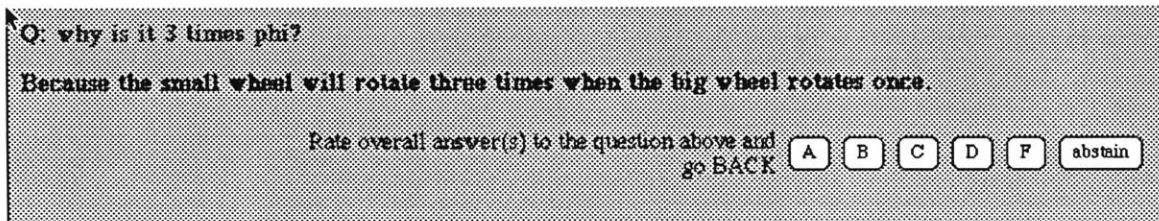
Figure 8: Section of an answer page with rating buttons

Instructor   Interface

Authoring Tools

The media management tools are grouped according the three media: audio, graphics and text. The audio and picture tools can only be used from a workstation that is running a web server. Figure 8 shows the author's view of an empty page.
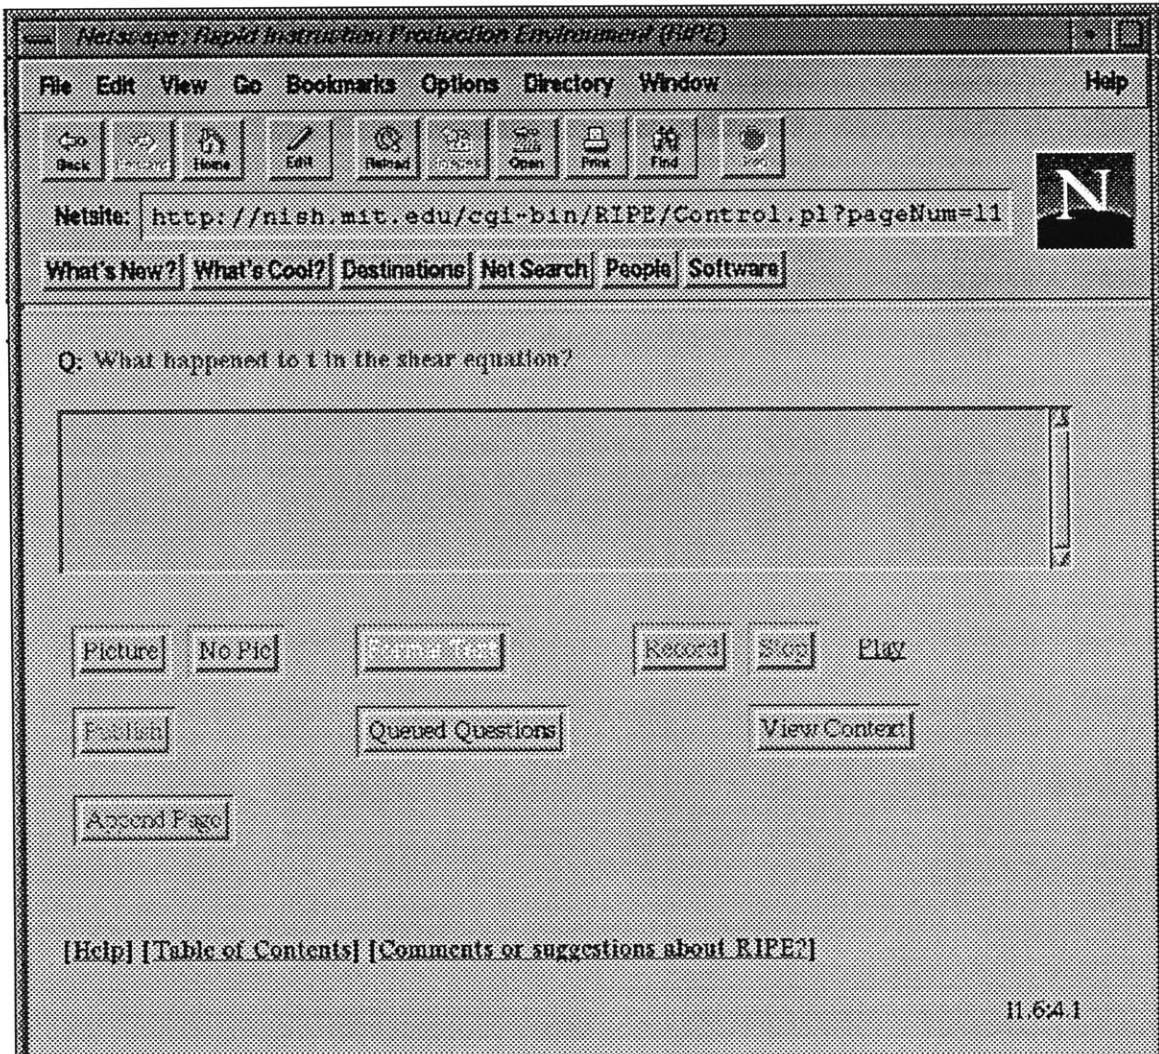
Figure 9: Author's view

Picture

There is one picture for every page. The first page of every answer automatically inherits the picture from the context of the question (ie. the page one level up). Two buttons, "Picture" and "No Pic" allow the author to take a new picture or erase the current picture. Any kind of freehand sketch, color drawing, or picture can be quickly captured once the camera is set up.

A special purpose photo-imaging platform, which includes a lighted stage, allows the user to set the size of the image by changing the height and focus of the camera. It is easier to do this while looking at a real-time image using the camera utility (see appendix with Instructor's Instructions) and this utility may also be needed to adjust the light level using the shutter speed.

The aspect ratio of the SGI camera is a little taller than an 8-1/2 by 11 inch piece of paper. If the height of the image is matched to the height of the page, this camera's resolution is such that drawings should be rendered with a felt tip pen. Higher resolutions are, of course, possible with better cameras.

Audio

There is one audio recording for every web page. The "Record" button initiates (within a half second or so) recording and "Stop" halts it. To re-record the audio, the user can hit "Record" again; the last recording will be replaced. Hitting "Record" followed immediately by "Stop" will effectively erase the recording.

Once a recording has been made, the link "Play" will appear. This invokes the RealAudio player with the recording. In the instructor mode, playback will not occur automatically when a page is opened.

Text

The author can type textual remarks into a box. Students will see this text field displayed as a caption below the picture. The text is actually interpreted as HTML but carriage returns are explicitly converted into HTML line breaks in order to preserve the paragraphs.

HTML commands may be used to change size, style and color of text if desired. The "Format Text" button allows the author to preview the output.

Publishing

When an answer, including one or more pages, is complete, the author should explicitly publish it. Publication will send out e-mail notifications to the interested students and remove the question from the list of queued questions. Once published, the pages will no longer be marked "under onstruction" and students will be able to ask questions.

Navigation

Paging

The instructor's "Next Page" and "Prev Page" controls work like the students' except that, on the last page, the "Next Page" button is replaced by an "Append Page" button which appends and goes to a new, blank page, the last in the linear sequence.

Queued Questions

The standing questions are listed on a page that can be accessed by hitting the "Queued Questions" button. The queued questions page (see figure 10) contains at least one line for each question. Each entry starts with a multi-part page number, followed by the a hypertext link - the question text - leading to a page where an answer can be composed. The last part of the each entry consists of the email addresses of students who wish to be notified when this question's answer is published.
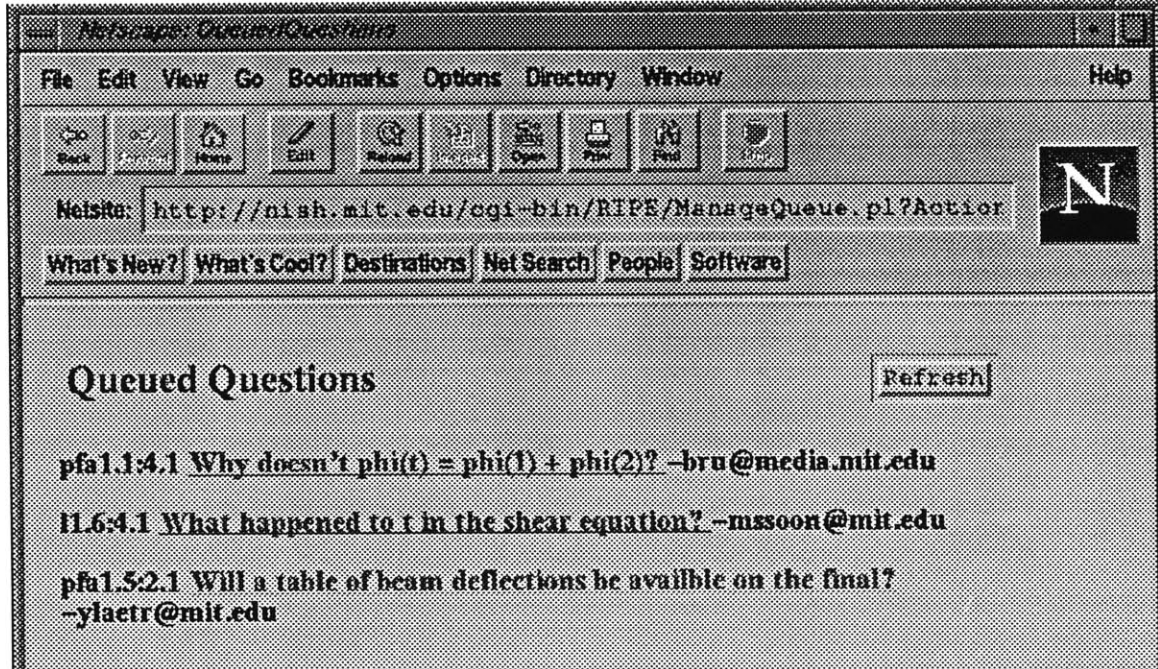
Figure 10: Instructor's view of questions waiting to be answered

Listed Questions

The author can also see the published and pending questions associated with each page. The view is the same as the students but, when the author activates a pending question, the link leads to a page with tools for authoring.

Pop on Publish

When an answer is published, the author is "popped up" to the page with the context of the question just answered. This allows the author to select any other pending questions from the same context while the memory is still fresh.

Pop to Question
Context

At any time while answering a question, the author can hit the "Context" button in order to see and hear the material which spawned the question.

Bookmarking

Both students and instructors may bookmark a page. This is presumed to be most useful to students who would like to interrupt their study and return later to the same page. The "Bookmark" button / text box generates a form that requests/confirms the user's e-mail address and, on confirmation, e-mails the text along with the URL of the page.

Help / Entering RIPE

There are separate help pages for students and for authors (Appendices A and B) Since there are some helpful instructions that students should see (particularly about not using the browser's 'Back' button to return to the context of the current page's question), the help page is also the preferred point of entry for students. The student help page then links to the presentations of interest. During the experiment, the student help page was referenced from the home page of [HITE].

A help link is also available at the bottom of RIPE pages. If the user is an author, the help button displays a different help page with detailed instructions for authors.

Table of Contents

A table of contents is automatically updated wherever the author creates or modifies a top-level page. The page number followed by the first sentence of the text field is used as the entry for the table of contents. The entry is also a link to the page itself. The table of contents may be accessed thorough a link which appears at the bottom of student and instructor RIPE pages.

User Authentication

Users must provide a username and password when they first attempt to access the authors' content. This authentication is provided by the WWW server. RIPE ran on a Netscape server which provides directory and file level access controls. At present, RIPE maintains all content in one directory and Common Gateway Interface (CGI) scripts in another. Appropriate (read or write) access to the authorized groups (students and authors) is allowed for these directories.

5 Trial

Testing proved to be the weakest link in this project. The ability to interact with teachers is the primary determinant of satisfaction with distance learning [Hiltz]. RIPE requires a faculty person to develop and refine the content and this is not a casual involvement. It should be no surprise that developing an interactive course from scratch will require at least as much effort as writing a textbook. (Such a large effort could and should be shared, of course.) Despite the investment required by teachers, there is currently little opportunity for remuneration.

The initial plan called for capturing a few lectures by recording the audio and either taking a picture of the blackboard with a digital camera or scanning a particularly neat student's notes. The plan was discussed with Prof. Michael Jordan for use with his class on Bayesian Networks. However, neither Prof. Jordan nor his TA could afford the time to handle the on-line questions. Dr. Sonwalkar also raised the possibility with mechanical engineering faculty but similar objections were raised. Furthermore, experience with HITE's mechanical engineering courseware strongly suggests that it would be inadvisable to offer RIPE as a option. Students are in the habit of going to class and aren't likely to make time for a supplementary on-line version covering the same material - even if it does hold out the possibility of asking questions. A collaborative approach like RIPE performs better if there are more users. If RIPE were to be offered as a merely optional component of a traditional class, it is unlikely that a critical mass of users would ever develop.

However, we were able to run a very limited test at the end of the semester.

Hypotheses

We have an interest in testing the following hypotheses:

(H0) **The number of questions per page will increase asymptotically with usage.**

This can be examined by plotting the number of questions against the number of uses for each page. A page may be defined to have been used if it has been rated because the act of rating implies some consideration and interpretation. However, RIPE's top-level pages are not rated. In that case, a page opening or "hit" may be equated with use. This could lead to misinterpretation when students are just casually leafing through the pages but this is probably not important for testing the idea that, in general, the rate of new question creation will decline and eventually become nearly zero.

H0 assumes that the audience remains relatively homogeneous. If the subject is JAVA and the initial audience is composed of C++ programmers whose new questions have declined, it can be expected that a fresh class filled with Visual Basic programmers would bring a resurgence of new questions. (In a case like this, it would probably be desirable to take background into account when ranking the questions See Futures section.) Another caveat for H0 is that something in the environment may have changed. e.g. A new airline crash, for example, could elicit a new and specific line of questions about a page on conditional probability.

(H1) **The average rating for a given question will rise as the number of its sub-questions rises.**

Additional questions will serve to clear up ambiguities in the answers, address related issues, and serve to fix answers that did not work well after the first attempt. If H1 is true, then H2 may also be true.

(H2) **Within a population of questions on any given page, there will a selection effect over time.**

Questions that are initially better rated will tend to attract more users. (Since the questions are ordered with the highest rated first, students are likely to be strongly influenced by the ratings when picking questions.) The more heavily used questions will tend to become more fully developed with follow-on questions and, if H1 is true, have higher ratings. Over time, we'll expect to see some divergence in ratings. The well rated answers will

receive heavy use and should improve while the poorly rated queries will receive less attention and improvement.

If there is a clear distinction between better and worse content, students can be expected to choose the better materials. Hence H2 could lead to H3.

(H3) The answers will improve over time.

The perceived quality of the content will depend upon the material that the user actually views. If the best Q&As improve and users tend to choose the best rated questions, then the average of all ratings will increase with time (and use).

Of course, the rating system is subjective and it is possible that students' standards may change, particularly if the first users are better prepared and less critical than the students who come later. The effects of early vs. late users within one class could be studied by running these test with multiple classes, allowing the questions to accrete, One important cause of paradoxical results may result from alienation due to too many questions being covered. This relates to the phenomenon cited by [Hiltz] that instructor contact is essential to student satisfaction (See Futures section for a discussion of remedies to this problem.)

Methodology

Prof. Frank Feng taught the Spring 1997 version of 2.001, Mechanics of Materials, and he agreed to participate in the trial. Dr. Feng provided three sets of source materials during the week leading up and the day immediately following the final exam: 1) Spring 1996 final exams with solutions, 2) Solutions to a sample final that had been handed out in class, 3) Solutions to the Spring 1997 final. All were scanned using Adobe Photoshop and cropped before being transferred onto web pages. The final exam with solutions was also rearranged so that the questions and solutions would appear on the same web page. The originals were not created with this use in mind and some of the smaller notations were difficult to read in their digitized form. There were no voice-overs.

Links were added to the 2.001 area in HITE and Prof. Feng told students in class to look for the new materials there. The class was also notified by e-mail. The author showed Prof. Feng how to use a browser on his desktop Macintosh to enter RIPE with instructor priviledges. Prof. Feng could reply

from his office to the questions that only required text based answers. The camera and microphone equipped RIPE server was a 1 minute walk downstairs.

No special instruction was given to the class. The experiment ran for approximately one week.

Results

The summary of use is presented in table 2:

	Pages	Hits	Qs	Q Hits	Ratings
Spring '96 Final w/ Solutions	10	2036	6	205	46
Practice Exam Solutions	6	882	6	103	13
Spring '97 Solutions	2	25	0	-	-

Table 2: Usage summary from the trial

Overall, there were only twelve questions and no sub-questions. Table 3 shows the text of the questions. Particularly for the Spring '96 sample, most questions were about the overall nature of the final rather than the material on the page. All but two of the questions were answered with a simple yes/no answer that did not lend itself to a rating. Therefore, it is not surprising that 78% of the users merely abstained or used the "Back" button instead of giving a rating[5].

Since almost all of the questions are irrelevent to the pages where they were raised and there are so few questions and no sub-questions, the data is not useful for an analysis of the hypotheses. The problem seems to be with the choice of content. The final exam solutions were self contained and, as Prof. Feng noted, if a student couldn't understand the solutions, it would be too late to pass the course.

On the positive side, no malfunctions were reported or are known to have occured during the trial.

---

[5] For the Spring '96 final, about half of the 46 ratings were "A", while the other half were "F". Presumably the "F" ratings had more to do with whether the student liked the answer than whether the reply answered the question.

Ratings	Avg. Rating	Student's Question
Sample Exam Spring '96, page 3		
12	3.16	Are there hard copies available? Netscape doen't like to print big jpegs
11	2.36	is multiple choice going to be the format of the exam?
5	0.40	Will the exam be heavily focused towards the end of the class, or will it be more comprehensive?
Sample Exam Spring '96, page 10		
1	5	Will nonuniform torsion be covered on the exam?
9	3	is the formula sheet from the final last spring available?
8	2.12	This entire format (not just the multiple choice) seems very different from our previous tests. Will the same be true of our final?
Practice Final Exam, page 1		
7	3.28	why is it 3 times phi?
0	-	In reference to #1, I thought that was accounted for when we found torque on the larger shaft.
0	-	our TA only had phi (t) = phi (1) + phi (2)
Practice Final Exam, page 4		
1	0	what happened to t in the shear equation?
Practice Final Exam, page 5		
4	4.25	Will a table of beam deflections be availble on the final?
Practice Final Exam, page 6		
1	3	Could you please suggest some problems from the text dealing with columns?

Table 3: Details of questions from the trial

Conclusions

The trial in this thesis served to test the software functionality, not the underlying ideas in this thesis. Any future test will require a much more careful selection of the instructor and also the material. Ideally, students will already be receiving instruction over the WWW in a class where many questions arise because of the nature of the subject material or the diversity of the audience. The created content will probably not have any financial value to the instructor yet the exercise will require substantial participation

from him. The instructor's effort can be reduced by allowing TAs and students to post replies. (See Futrher Work)

6 Viability

Where is a RIPE like system most likely to succeed? Does it fit with the traditional or the open university model today? in 2010? What are the other applications? Is RIPE technically viable?

When Should RIPE be Used?

Content Newness, Student Heterogeneity (and Distance)

RIPE is really a methodology for rapidly creating content. Figure 11 shows two factors that, in addition to distance, favor RIPE:
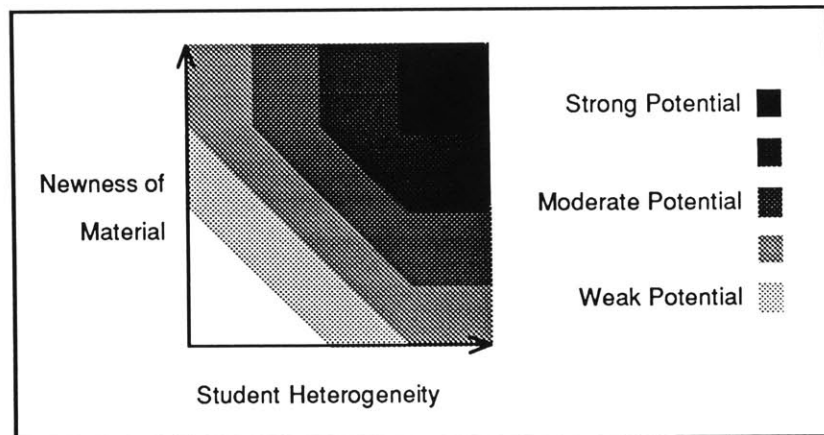
Figure 11: Factors influencing RIPE's penetration potential

The factors suggest that when the students come from a relatively uniform background (by way of admissions requirements or course prerequisites), have similar objectives (get a good grade; graduate) and are taking

fundamental courses that are invariant from one semester to the next (the instructor can predict the vast majority of the questions), RIPE offers the fewest advantages. The existing materials are already well matched to the typical student's needs.

Accordingly, it is hardest to make the case for RIPE in a core subject like mechanics of materials that, according to Prof. Feng, "has been taught the same way for a hundred years" to uniformly prepared students who are already on campus and expect to take classes.

On the other hand, if the subject is new and changing then there will be a need to develop content. Similarly, more heterogeneous classes will increase the need for content creation. In continuing education and graduate studies, students are more likely to come with specific issues, more diverse backgrounds and a greater variety of questions.

The third major factor is distance which, obviously, makes all of the on-line approaches more attractive.

Influence of Institutional Factors

Many of the most vigorous adopters of distance learning have been second and third tier schools. Usually, these are regional schools with commuting students. Their products are highly commoditized so proximity is the main reason that students select them. Since these will probably be the first schools to face serious revenue shortfalls due to rising enrollment in less expensive distance learning options, we should not be surprised that so many are responding proactively by moving towards cyberspace. So far, however, the content is indigineous and this limits the benefits.

In entertainment, 2nd rate local performances can no longer draw a crowd when the alternative is to see and hear nationally and internationally acclaimed artists at the movies, on CD and TV. Yet this "star performer" effect is little noticed in education. There are thousands of mediocre physics high school teachers who kill off intrinsic curiosity in tens of thousands of students every year. This, despite the fact that outstanding videotapes covering the same material have existed for some time. Why is the diffusion process so slow?

Technological diffusion can be held up by "gatekeepers". In this case, the teachers are threatened by the innovation and, since they are also gatekeepers, they have the power to block it. A mediocre accounting 101

teacher at a community college doesn't want his students to see how much better videotaped instruction could be. He is reluctant to give-up his role as the leader of the class and unsure of how he'd redefine his job[6].

Thus, established colleges will face a dilemma. If they promote the use of courses developed by star teachers at the expense of their traditional classes, they will alienate faculty whose role is to "stand and deliver". Yet, as on-line learning matures, its products may begin to deliver better value and students will gravitate towards programs that use them. Unless the non-star faculty are able to provide value in some other way, they will create operating losses. Newer institutions developed around distance learning approaches will not have the same high costs sunken into constructed classrooms and tenured faculty. This will provide them with an opportunity to compete in the top ranks to grant degrees that do not require residency for other reasons, such as research.

Though distance learning can be a way to leverage the faculty, one paradoxical effect has been that "the (MIT) departments with the lowest enrollments (e.g. Aero/Astro) have been the most agressive adopters of distance learning" [Lerman]. Whereas these departments hope to see revenue enhancement, heavily saturated departments at MIT see distance learning as a magnet for additional students that they do not want. Stanford has inverted this logic and turned its high demand courses (EECS) into sources of revenue with exceptional profit margins. The money comes from the (employers of) students who pay full tuition but place few extra demands on the campus infrastructure and faculty.

What are students really buying? Is it learning or is it the degree and the reputation of the granting institution which counts most? Is a reputation made by offering courses that are relevent to the real world, or programs that create a kind of obstacle course which eliminate the weakest students and toughen those that eventually graduate? If the latter is true, then RIPE's central objective could be seen as irrelevent but since selection of

---

[6] The job anxiety is not altogether justified. Stand-up teachers can become coaches and mentors to students. It may be desirable to delete some of the content created by RIPE, just to let students "discover" the missing answers with guidance from other students and the teacher. If all of the questions and answers are too clearly in place, the student may become unmotivated. This is part of the syndrome of "drill and practice" computer aided instruction, which was wisely derided by Prof. Seymour Papert as, "the computer programming the student (rather than the other way around)".

graduates by attrition is no longer a primary institutional objective[7]. In any case, RIPE should succeed in elevating the level of the students' knowledge and harder challeges can always be created.

Interdisciplinary studies are becoming increasingly important yet it is often difficult to cross-over into other fields because of specifc assumptions about the prepartion of the students. Prerequisites are usually required in order to ensure uniform preparation for the class but it is often impossible for students from other fields to take so many subjects. Usually, only portions of the prerequisites are actually needed for the higher level class. One approach is to break classes down into smaller modules so that the prerequisite portions could be isolated and taken individually on an "as needed" basis. RIPE can also help to fill in the prerequites for the students who need it, when they need it. This micro-modularization of the programs and the courses will be nescesary in order to accomodate more interdisciplinary students and participation from non-matriculating persons, including alumni[8].

The idea of strong alumni participation is attractive because it will help to align the curriculum with the prevailing needs. Input from working graduates also gives students a better understanding of their career options and helps to motivate their studies.

Metcalfe's Law

"The value of a network is proportional to the square of the number of users."

The more users there are, the larger (and better) the question database. If there are many users, it will be also be easier to quickly find someone to help. Both the quality and the speed of reply are expected to be key determinants of student satisfaction so, in general, large classes will be preferable to small ones.

---

[7] Though, up until the late 1960s , one-ninth of MIT's freshman class was flunked out. In the early 1980s, when the EECS department was oversubscribed, the entry level courses purposefully flunked a high proportion of students in order to discourage them from majoring in the program.

[8] Prof. Ricard Larson, in suggests the idea of MIT lifelong enrollment whereby alumni can participate in MIT courses for life, primarily via distance learning. [Larson]

Trickle Down from Corporate Desktops?

Teaching systems, which distribute the knowledge of experts, are the first step towards modelling and supporting many other kinds communications that are used in the real world. While RIPE's roots are in education, there is a much larger market for communication within and between corporations. Given the economies of scale in software development, it seems quite likely that systems designed for the broader market will trickle down to education rather than the educational systems being adopted and expanded for broader uses. Microsoft's Powerpoint, for example, already incorporates audio and is capable of generating presentations on the web.

Short Cases for/against RIPE

The following provides some examples of schools or programs and an analysis of the suitability of using RIPE in each case.

Open University

The Open University [Open University], based in Milton Keynes, United Kingdom, is the largest of the worlds' distance learning universities with 2.5 million graduates. The enrollment is 160,000 with 16,000 studying from other countries. Each student is assigned a tutor and most of the teaching is accomplished via video, texts and workbooks. There are 13 regional centers around the UK where classes are held. The Open University is beginning to offer a handful of courses via the internet and have established their own points of presence (POPs) in the UK. They plan to offer optional internet audio in the near future.

To the extent that the courses are on well understood subjects that have been taught for a long time, RIPE is unlikely to play an important role at the Open Universtiy. On the other hand, even though the proportion of candidate courses may be relatively small, this is a very large school which includes at least 10,000 graduate students who are learning at a distance.

HBS

The Harvard Business School has recently invested $11 million in IT infrastructure including custom software which comprehensively supports classroom instruction with on-line class readings, schedules and activity announcements. Regular threaded discussion groups are also supported. (Additions to the discussion are prominently announced to students, a feature which is designed to help keep the discussions active.) [HBS]

The software developed for HBS is intended to support, but not in any way replace, face-to-face classroom interactions. It is probably not in HBS' best intests to create an integral distance learning program. To a large extent, the value of the business degree comes from the personal acquaintances formed by the students and the network of alumni. At least for the forseeable future, distance learning will not offer as rich opportunities for personal interaction, and offering the program on a larger scale would dilute its exclusivity.

NJIT

The New Jersey Institute of Technology has been a pioneer in on-line courses. It now offers about 80 courses via loaned videotapes and on-line computer conferencing / e-mail. Approximately one-quarter of NJIT's current enrollment is from remote students (300 graduate; 200 undergraduate). [NJIT]
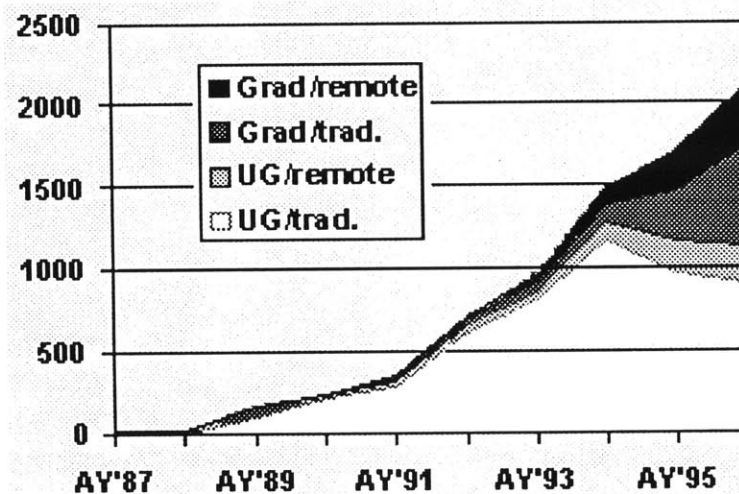
Figure 12: Enrollment Growth at NJIT

NJIT, like the Open University, is relatively experienced with distance learning. However, NJIT is much smaller and teaches most students on campus. It does not have the same economies of scale as the Open University and will sooner face pressures to integrate content from outside the school. Still, a RIPE like approach could be used by local faculty and students to hold discussions and fill in details of courses that are usually produced elsewhere. In a few cases, there will be star teachers among the faculty whose on-line 'lectures' and interactions with students could become the content which other schools integrate.

Stanford

Since 1969, the Stanford Instructional Television Network (SITN) has distributed courses from the School of Engineering to corporate sites. [Stanford] Under the auspices of the Stanford Center for Professional Development (SCPD), Tutored Videotape Instruction (TVI), broadcast and compressed video are also being used to deliver advanced engineering short courses, executive education, degree and non-degree programs. Applicants to the degree program compete with regularly enrolled students for graduate degrees. In addition, professionals may enhance or change career options by enrolling (withough matriculating) in one of the certificate programs that exist in Aeronautics and Astronautics, Computer Science, Electrical Engineering, and Engineering-Economic Systems & Operations Research.

Stanford is equipped with several teleclassrooms and a link to a microwave station which covers the San Francisco Bay Area. There is a broadcast schedule and remote sites may participate by phoning in. Instruction by videotape was concieved as a way to reach students beyond the broadcast range and internationally. However the mode of use is differentiated from broadcast by the fact that an on-site tutor, chosen by the participating company, plays a much more interactive role. He may stop the tape to initiate class discussion, and answer any questions that may be specific to the company audience. However, direct broadcasts via microwave are the prevailing mode of dissemination.

For the last five years, video-conferencing has also been used in conjunction with fax, email and overnight delivery. It has been uniquely valuable for "advising sessions, tutorials, office hours, and to support faculty/industry communications in consultative and research-related opportunities."

SCPD has recently begun to experiment with on-demand formats via the web. Downloadable content in QuickTime and Xanim are available for a half dozen courses. Most recently, SCPD has begun to offer streaming video using the latest technology from VXTreme.

The Stanford case contains a spectrum of offerings from, at the lower end of the range, mass communication of more fundamental courses via one-way media (microwave TV network), to slightly more customized courses (videotape with a live tutor present), to presentations that are highly interactive and tailored specifically for an individual client (videoconferencing).

Stanford has star talent since the instructors are often also leaders in their fields. While they may not be the best teachers, they may be regarded as the best qualified authorities on what to teach. Therefore, they could set the syllabi and provide the top level of the material, leaving others who are more in touch with the students to fill in the gaps. RIPE is designed to do this. The biggest barrier to acceptance may be that SCPD and its corporate clients are tied to the video medium. Distribution should not be a problem (the companies are well networked and the students will almost always have a PC on their desks) but it will be difficult to convince the personnel to shift over to a new technology. Still, Stanford has most of the right ingredients (cutting edge courses and graduate students who are already distance learners) and could be an attractive place to run trials and develop this technology.

Gartner

The Gartner Group [Gartner] offers over 300 courses related to computer application use, systems administration and programming in multiple delivery formats including CBT (computer-based training), multimedia CD-ROM, video, and the Internet.

Gartner is one of many groups that offer training on job specific tools, mostly software applications for office use. FAQ compilation in the context of the training materials may be useful, especially if the audience has highly varied levels of education. However, in general, the material does not have very much depth and this may mean that almost all of the questions can be anticipated by the course designers. There is less value in having an inexpensive and rapid content development system since the courses are sold in volume and the per unit authoring costs are low.

Technical Viability

Reliability, Performance, Reusability, Security and Statistics

There are no serious technical barriers to implementation. There is room for some unreliabiltiy since the application does not have to be running 24 hours a day. Users of the web are used to, and tolerant of, temporarily inaccessible sites.

Web server performance will degrade gracefully under load and the system can be distributed across multiple server (see below). A single desktop server (SGI Indy) can serve up to a few thousand interned users while a larger dedicated web server (SGI Challenge) can serve up to 10,000 to 50,000 users of a single course. The hardware is becoming increasingly inexpensive as the new multi-CPU pentium pro systems can also handle about 100,000 to 500,000 hits per day given adequate network bandwidth. A T1 line (1.544 Mbps) will only carry approximately 10,000 hits per day so bandwidth is actually the limiting resource.

The content should be completely reusable except that translators will have to be supplied which convert older data formats.

Security concerns are the same as for other web sites running CGI scripts. It is important to set the proper permissions and to ensure that the scripts themselves strip out any commands which could, for example, directly access the server's OS.

Gathering statistics is intrinsic to RIPE. In addition to keeping a running count of page uses and average ratings, RIPE maintains a log which can be used, ex post facto, to reconstruct desired information. Currently, the log tracks data about page accesses, ratings, and page publication.

Scalability

Server

RIPE's single server design has limits of scale. If it became nescessary to expand the user population beyond the capacity of a single server, the database could be copied and content development could proceed on parallel tracks. In some cases, it might actually be desirable to start independent content development tracks in order to neutralize the first mover advantages. (See the discussion below under Future Work.) However, it will eventually be nescessary to synchronize the databases so that the best creative efforts can be collected and propagated. RIPE currently identifies web pages using the following addressing:

<presentation id>.<page>:(<sequential question number>.<page>)

This addressing convention will have to be modified in order to synchronize web pages on multiple sites. To prevent conflicts arising from duplicating the <sequential question number>, a site ID could be prepended.

Database synchronization may be planned by a central authority (as in Lotus Notes) or allowed to develop in a distributed way (as in Usenet). Central planning can guarantee that one server will have the most up-to-date database and this is likely to be the first step in scaling. This approach can also keep the central database administrator up-to-date on the usage at each of the satellite servers - which may be important for course administration and billing (until internet micro-payment arrives).

A hierarchical server arrangement will work well for distribution between universities. Top schools could own and maintain the high level content with contributions coming back up from the 2nd and lower tier schools. The top-level school would also process all of the content updates and be in a position to edit and stop errors from propagating.

The database represents a threaded discussion like Usenet (internet news groups) except that the items are rated. Usenet's peer-to-peer update strategy could also be applied to the database update problem. Internet micro-payments will help to promote this approach. But, since there is no central control, it won't be possible to exert the same editorial oversight.

The ratings are likely to become a significant, if not predominant, part of the traffic needed for updates. Ratings can be summarized to reduce the bandwidth.

Authoring Client

The hardware costs of the mini-studio are in the $300 - $1000 range and should not pose a significant problem for adopting RIPE on a wide scale. At the low end, Connectix offers a color 640x480 camera and interface card for use with PC and Macintosh computers for $300. Elmo is the prevailing in the destop videoconferencing marketplace and offers a stand and camera for about $800.

Scanners may be used in lieu of a camera and, although they are less convenient, they do offer superior image quality.

7 Suggestions for Further Work

Pre-trial Updates

The following changes can be quickly made and will likely prove to be quite useful. It is assumed that the trial instructor(s) will author using a machine which also serves as a web server.

Multiple Authors

The restriction that only the instructor can answer questions increases a key determinant of student dissatisfaction, Mean Time Til Answer (MTTA), and may also create an unacceptable workload for the sole authorized teacher. The required design changes are minor: 1) enqueued questions must be locked while one teacher is working on the answer (the lock may be relinquished to allow someone else to complete the answer; published questions would disappear from the queue as they do now) 2) answers should be tagged with the authors' names since students may develop preferences (and people like to see their names next to high ratings!).

Student Replies via Text

While this is technically simple, it will probably be essential to user satisfaction, especially in post-graduate and more dynamic fields where RIPE has the greatest value. While allowing students to help eachother will improve MTTA, their interactions might also devolve towards communications that are more typical of "chat" rooms. Particularly when students are the creators of content, they may not be the best arbiters of its quality. The next two sub-sections suggest specific remedies to this problem.

The implementation is a straightforward follow-on from supporting multiple authors. Students must also have access to the queued questions and be

permitted to publish replies.

"Seal of Approval"

Once the non- experts are allowed to become authors, the lliklihood of incorrect or overly simplistic answers increases. The answers should be officially approved before they become permanent additions. This will require the creation of a new queue consisting of questions that have been published by uncertified users such as students and TAs. There could be several options:

1. authenticate
2. discard
3. return with comments for revision
4. edit and authenticate

In any case where changes are required, it should be possible to notify students who have already seen the mistake. This may not be exercised by default but should certainly be available in instances where the error is egregious. (This is the most essential of the set of activity notification mechanisms described below.)

Authenticated answers will bear a distinguishing seal of approval, like a green check mark, next to the corresponding question.

Group Similar Questions

The ratings are useful when the user is selecting between similar questions. Therefore, like questions should be grouped together. Students who are dissatisfied with the existing answer will be the most inclined to repeat and, perhaps, rephrase the question. When a poor rating is given, there should be a dialog which allows the rater to: 1) repeat or rephrase the question and, 2) state why the previous answer was unsatisfactory. The instructor should be able to see the student's complaint in context before attempting another response. The new question should be grouped with the previous one(s)..

Import Images

Although few students can be expected to have cameras or digital tablets, some will have the inclination and time to use other programs to add graphics to their answers. Uploading and including standard image files (.gif .jpeg) should be supported.

Change the Rating Labels

Some faculty members will be offended by the idea of students assigning grades to their explanations. The "A" through "F" scheme does not soothe these sensibilities. More neutral labels should reflect the dimension being measured. (e.g. was the question answered? vs. was this useful to you?)

Rank Questions by Hits Received

Currently, well rated questions, though they may be of limited interest, appear at the top of the list. Re-ranking the questions according to the number of hits they've received will make it easier to locate the most useful ones.

However, rank ordering by total number of users introduces a strong bias towards the older content. In order to reduce this bias, it would be better to rank using the more recently collected data. There is a tradeoff in deciding where to set the cut-off between too old (biased) and too recent (not enough data to overcome sampling errors).

Groups of similar questions (see above) should remain grouped together. The group should be ordered according to the amount of traffic received by its most popular question. However, within the group, the questions are direct competitors and should be ranked by average ratings. This will give new Q&As a better chance of being selected and improved. (There is again the problem of deciding the period over which to average the data.)

Short Term

Picture and Audio Capture from Web Clients

Currently, multimedia content can only be captured on a web server. Since the vast majority of instructors' office machines cannot run a web server (they're not multitasking and may not have adequate performance), instructors would have to go to wherever the server is installed. This inconvenience increases the MTTA and decreases student satisfaction. Client software should be developed to capture images from cameras and microphones attached to users' PCs.

Students should also have access to these multimedia input facilities. Some will become successful as self-appointed tutors, especially if they are remunerated.

Position Questions on Top of Images

Many questions will pertain to a particular point (or area) of the image. Rather than forcing questioners to use text to describe what they are referring to, which can be quite cumbersome, the entire graphic should be clickable, leading to a dialog where the question's text can be captured.

There are several options for displaying question location. One would be to overlay a transparent question icon over the image which could be clicked to show the question's text. However, it may be too distracting to show the locations of all the questions by default. Clicking in the image could show the questions, clicking again on a quesiton could bring its text to the foreground, or top of the list. Variations on the transparency and size of the question icon could convey information about its rating and usage.

Medium Term

Re-implementation using JAVA

The following features require more precise control of the client than HTML and JavaScript allow. In order to implement them, the server would start the session by downloading a single web page containing JAVA code that takes over the client/server interactions and the user interface.

Pre-Fetch along Most Likely Paths

Delays, experienced while waiting for images to download and audio to start playing, waste students' time and detract from the overall value of web browsing, let alone courses authored using RIPE. This can be minimized, to some extent, by pre-fetching pages. It won't be possible to pre-fetch every possible selection but the data collected on the popularity of each of the various options can be used to boost the probability of sucessful pre-fetches.

Continuity of audio can engage the user's attention while pages are being downloaded. In the event that the correct page is not pre-fetched, the audio should be started immediately. (The only exception to this would be for clients who must connect at about 30kpbs by modem. While the audio is being transmitted, there is only about 10kbps left over for the image. A 60Kbyte .gif file would take approximately one minute to download.)

Animation

Very simple animations, such as the movement of a pointer or highlighting particular regions of the image can, when synchronized with voice, serve to focus attention and eliminate the need for descriptive verbiage. Typically, the instructor will capture an image to screen. Then, while recording or playing the voice-over, he could use his mouse to show the areas of interest. The bandwidth requirements for this kind of animation are minimal.

Synchronous Communication

So far, the RIPE transactions have been asynchronous, like e-mail. Assuming that the question queue is well tended, a reply can be expected reasonably quickly. But the exchange is very different from opening a direct audio channel and creating an electronically shared desktop which both parties can see. Full video is not nescessary. Frequent image updates and continuous pointer (mouse) tracking from both sides should be effective in conveying the feeling of a shared desktop - even if only the instructor can draw.

The synchronous option is quicker, creates a sense of immediacy that comes from working live, and may allow the student to help the teacher to focus the reply. The students' sense of satisfaction may actually be increased by hiding some of the previous questions and answers, leaving these instead as topics for exploration, especially among students themselves - even if they have no cameras, or even, microphones.

Asynchronous communiction will still occur and the teacher may prefer to keep some interaction asynchronous because it allows more time to organize a reply and present it clearly. Alternately, they may choose to interact in real-time at first and then go back to improve the recording.

Place Questions in Time

Questions may also arise at particular moments in the audio track. The best way to capture and represent these questions should be a subject for further research. Should questions with temporal links be rated and ranked by use like the other questions? Or should they gain visibility at the moment when they were asked? Or, like the physical references, should they remain invisible (except perhaps for aural cues in the soundtrack) until a student attempts to raise a question?

This feature should be implemented together with synchronous animation. Both will require interfaces with the audio player and data with temporal linkages.

Error Flagging by Students

This extends the functionality outlined above, wherein content experts have the authority to flag errors and bring them to the author's attention. This feature would allow students to flag suspected errors which are then brought to the special attention of the instructor.

Activity Notification

As long as the content is developing, students who are early users will miss the developments that follow. This is especially critical if errors are found in the materials which the early users have seen. Students may also want to be notified about questions that follow-on to their own, or any questions which they have not viewed whcih have become highly popular.

Credit Allocation

Some of the best content creators will be students. The efforts which result in enduring contributions to the corpus they should receive some form of credit. One way to account for the credits would be to measure total hits on the author's contributions. Students might also get credit for contributing questions. Credits could count towards grades or have cash value, like royalty payments. The ability to make micro-payments across the internet will be an enabling factor for the latter.

Long Term

Semantic Analysis of Questions

Students may type in the text of their questions and a semantic analysis can be used to match the entry agaist the recorded questions. Relatively shallow analyses should yield useful results and can be more robust with respect to syntactic variations than deeper analyses [Whitehead]. There are several reasons to have students type questions:

1. The interaction with the system will more closely approximate interactions with a tutor (the original goal). If the question cannot be

answered out of the database, the system can then engage a human to help.

2. Once all of the common questions have been recorded and listed, the student is in danger of glancing at the list without ever really allowing his own thoughts to form fully. Forcing the student to actually type in his own question may give him the opportunity to think for himself.

3. Although an answer may be very good, the corresponding question may not be as well stated as it could be. Questions could simply be restated and grouped manually so that the questions will compete to represent a particular answer (rather than several answers competing to respond to a given question). On the other hand, semantic analysis could help with the task of grouping questions. However, when students rate an answer based on whether it answered their question, they may in fact be rating the analyzer rather than the answer itself.

4. In the event that the question database has grown very large, it may be faster to type in a question, or at least a key word to help narrow the search.

Peer Clustering and Bayesian Analysis

Another way to predict the questions that a student will want to see is to gather a history of the questions that he has selected (and rated highly). The history can have several components and may be used in several ways:

1. The pattern of questions asked can be used to identify other students who tend to ask the same questions. These students form a peer group and the ratings of the students who have gone ahead in the course will help to predict the preferences of the students who follow.

2. Using bayesian analysis, trajectories through the material can be analyzed. e.g. This would allow the server to calculate the probability that the student would want to see question 4 on the next page given that he had already activated questions 2 and 6.

Multiple Rating Axes

The above pattern analysis techniques could be boosted by collecting more data. In addition to rating whether the given question was answered, students might be asked to rate their feelings about whether the answer was satisfying, useful or interesting.

Eventually, this might be combined with work on affective computing. By monitoring emotional state, via sensors heart rate and galvanic skin

response for example, it may be possible to implicitly determine the
effectiveness of media being viewed.

Conclusions

Video and computer mediated communication, are relatively recent inventions which are significant because they can sharply increase the accessibiltiy of learning opportunities. Historically, significant technological advances, though intially cobbled together with existing processes and organizations, eventually develop an infrastructure that supports an entirely new modus operandi. In this case, we should expect to see new models for educational institutions and the interactions between teachers and students within them.

RIPE supports a methodology for creating WWW content that is significantly different from the first wave of adapted processes. (1) Rather than using WWW programming specialists to create web pages, RIPE would provide the content expert with a "mini-studio" which he can use to quickly create and publish new instructional approaches. (2) The process, unlike the WWW programmer reliant approach, allows new inteactive content to be created in an interactive setting rather than republishing material that was originally destined for a printing press or a videotape. (3) Students use the new materials soon after publication and provide feedback in the form of explicit ratings, as well as follow-up questions. (4) The focus of the RIPE enabled development efforts is on creating and choosing content for what it says rather than for how it looks. The material can be imbued with higher production values once the most useful approaches have been identified.

While none of the proposed technical enhacements to RIPE requires a breakthrough technology (though this does not preclude the possibility of discovering significant approaches during the implemenation), there are numerous non-technical barriers. (1) There are large investments in content using the old media. (2) Most educational institutions are built around classroom teaching, with which RIPE would compete. (3) Faculty do not know how they would redefine their roles if their lectures and recitations were made obsolete by interactive materials from star performers.

These objections suggest that the optimal conditions for early adoption will be: (1) New subjects, or at least ones that are being taught in a different way or to new audiences. (2) Classes already using distance learning. (3) Faculty who are highly respected experts or star performers with broad appeal (i.e. the displacers rather than those who fear that they'd be displaced).

The most difficult part of this thesis was finding a faculty member willing to use RIPE in their course.

The work undertaken in this thesis has shown:

1. the concept of a "mini-studio" that allows nearly instantaneous content creation for the World Wide Web using the sketch and talk methods with which teachers are familiar.

2. an implemenation based on server side PERL scripts

3. a trial which showed no significant problems due to loading or software errors

4. that recruiting teachers to use this new approach is difficult because RIPE competes with the classroom approach and teachers currently have few incentives to use RIPE.

5. that the type of material used with RIPE should be carefully chosen because some materials do not elicit many questions.

6. that the proposed rating system does not work for some kinds of answers (i.e. yes/no replies to administrative questions)

The work suggests that teachers, at least those who are already involved in on-line classes, will use mini-studios to create content themselves, often in synchronous interactions with students. It is also likely that students will have access to mini-studios and use them to create materials of their own. The pletora of resulting output will need to be managed and this thesis suggests that ratings, as well as oversight by content experts, will play a role in helping to select among competing answers.

It is planned that the Hypermedia Teaching Facility will continue with further testing. RIPE has been largely integrated into the HTF's HITE course delivery platform. A broader set of course offerings is important for attracting end-users and HITE already offers several courses as well as a framework for administration.

References

[Ackerman 96]   CSCW '96 Proceedings Answer Garden 2: Merging Organizational Memory with
                  Collaborative Help Mark S. Ackerman, David W. McDonald

[Ackerman email] From private correspondence with Mark Ackerman on 10 April 97
                  <ackerman@sima.ICS.UCI.EDU>

[Adobe]   http://www.adobe.com/prodindex/pagemill/main.html

[Avalon]   Avalon Information Technologies, Inc., Brampton, Ontario, Canada
                  http://www.atlantis.com/~avalon/

[Centra]   Centra Software, Inc., Lexington, Massachusetts http://www.centra.com/

[CityU]   http://kcox.cityu.edu.hk/papers/et2000.htm

[Cyberprof]   http://cyber.ccsr.uiuc.edu/cyberprof/general/homepage/Newpage/first.html

[Deming 94]   The New Economics for Industry, Government, Education, W. Edwards Deming MIT
                  CAES 1994

[Farinetti]   Farinetti, Laura & Malnati, Giovanni. 1996. Remote Tutoring: What We Learned by
                  a Practical Experience. Educational Telecommunications, 1996, Proceedings of ED-
                  TELECOM 96 Boston, Mass., USA. 97 - 102.

[FirstClass]   http://www.softarc.com/

[Gartner]   http://www.gartner.com/training/products/alpha.html

[Gloor]   Gloor, Peter A.; Makedon, Fillia; van Ligten, Oliver. 1996 Obstacles in Web
                  Multimedia Publishing; Bringing Conference Proceedings On-Line. Educational
                  Telecommunications, 1996, Proceedings of ED-TELECOM 96 Boston, Mass., USA.
                  121 - 126.

[HBS]   http://www.atg.com/dynamo/press/profiles/hbs.dyn

[HITE]   Statistical Analysis of Hypermedia Teaching Facility Usage, Spring 1997, HTF
                  Publication available from Dr. N. Sonwalkar, MIT Rm 9-230,
                  http://curricula.mit.edu

[iFactory]   Interactive Factory, Inc., Boston, Massachusetts http://www.ifactory.com/

[Larson]   MIT Learning Networks (Draft), Richard C. Larson, Director MIT CAES

[Lerman]   Comment made during a meeting on May 22, 1996

[JALN]   Journal of Asynchronous Learning Networks, http://www.aln.org/

[Macromedia]   http://www.macromedia.com

[NJIT]   http://www.njit.edu/DL/sld002.htm

[Open University]   http://www.open.ac.uk

[Rattner 96]   Collaborating with Human Factors When Designing an Electronic Textbook,
                     LETT'96 (Leading Edge Training Technologies Conference) Julie Rattner

[Shiba 93]   A New American TQM: Four Practical Revolutions in Management, Shoji Shiba,
                     Alan Graham, David Walden 1993 Center for Quality Management

[Stanford]   http://scpd.stanford.edu/intro.html

[Stifelman]   Proceedings CHI '96. Vancouver, Canada Augmenting Real-World Objects:A Paper-
                     Based Audio Notebook http://lisa.www.media.mit.edu/people/lisa/

[Suppes]   http://w3.scale.uiuc.edu/education/trustee/Stanford.html

[UIUC CyberProf] http://cyber.ccsr.uiuc.edu/cyberprof/general/homepage/Newpage/toplevel/welcome.html

[Utterback]   Innovation, Competion and Industry Structure, International Center for Research on
                     Management of Technology, 1991 Working Paper #29-90

[Webcrossing]   http://webx.lundeen.com/

[WebCT]   http://homebrew.cs.ubc.ca/webct/webct.html

[Whitehead '94]   AutoFAQ, GTE Laboratories working paper, Steven Whitehead 1994

[Whittaker '93]   Analyzing the intonational contour of natural speech. Technical report, HP Labs,
                     1993

[Whittaker '94]   CSCW '94 Proceedings Filochat: Handwritten Notes Provide Access to Recorded
                     Conversations Steve Whittaker, Patrick Hyland, Myrtle Wiley

Appendix A: Student Help Page

Figure 13: Student Help Page (top)

Help

Do NOT use the 'Back' button on your browser. All navigation within RIPE should be performed using the buttons on the page. Use 'Next Page' and 'Prev Page' to move forward and back within an explanation.

Questions take you down one level to an answer. In order to go back up one level, rate each answer using the buttons ('A', 'B', 'C', 'D', 'F') provided. When rating an answer, please judge how well the stated question was answered. If you visit a question that doesn't interest you, you can 'Abstain' in order to go back to where you were. The instructor will not use your vote.

The numbers in the left column next to the listed questions are their average ratings (5 point scale).

If you have a question that is not already listed, ask it! You may also repeat a question if the existing answer is inadequate. (Note: Your questions are not automatically signed; if you want, you can sign them yourself.) Jokes are encouraged.

Other Notes

The 'Bookmark' button will send you an email containing your comment and the URL of the current page. Use the bookmark when you suspend your session, to easily find where you left off.

These pages may contain audio. To use your own headphones on SGI Indys, reach behind the right hand side, the first thing you'll feel sticking out is the headphone jack. On Suns, the jack is in the middle of the back, you have to go behind to see it.

Note: Under some circumstances, use of the 'Back' button will cause errors as well as defeat the purpose of rating answers. To help prevent accidental use of the 'Back' button, go to Netscape's 'Options' menu and select 'Show Toolbar'. Thank you for giving us your feedback.

• •

If you have comments or suggestions, please email me at _____

Figure 14: Student Help Page (bottom)

Appendix B: Instructor Help Page

Instructor's Operating Instructions

Instructor's Operating Instructions

The following steps will tell you how to create your presentation, and develop it, using feedback from students.

1. Go to the machine nish.mit.edu and log in.
2. Open a Netscape browser. (Type "Netscape" into Desktop: UNIX Shell.)
3. Open the HITE home page (URL: http://nish.mit.edu)
4. Select "Other Projects". Select "R.I.P.E".
5. At the bottom of the page, select a presentation.
6. Enter your username and password.
7. You will see the table of contents for the presentation. Select the first entry, or page.
8. You are now seeing the instructor's view of the presentation. At the bottom of the web page, you will find a link to Help. You may want to follow this now in order to get the latest version of this document.
9. Turn on the lights on the camera platform.
10. Double click on the camera icon. Open the Tools menu; select Video Panel. Open the Pro menu; select Indy Cam: Signal Controls. Set the shutter speed to 1/125. This should give the best image quality.
11. Important: close all of the camera controls that you just opened.
12. Put the battery in the microphone. From the desktop palette, select Desktop: Audio Control. Click the Meter button. Talk into the microphone so that the meter peaks in the yellow region. This the way that you should talk when recording audio for your presentation.
13. Turn up the Speaker on the Audio Panel.
14. Stretch the top and bottom of the browser window to fill the monitor vertically. (This will make it easier to use the controls.)
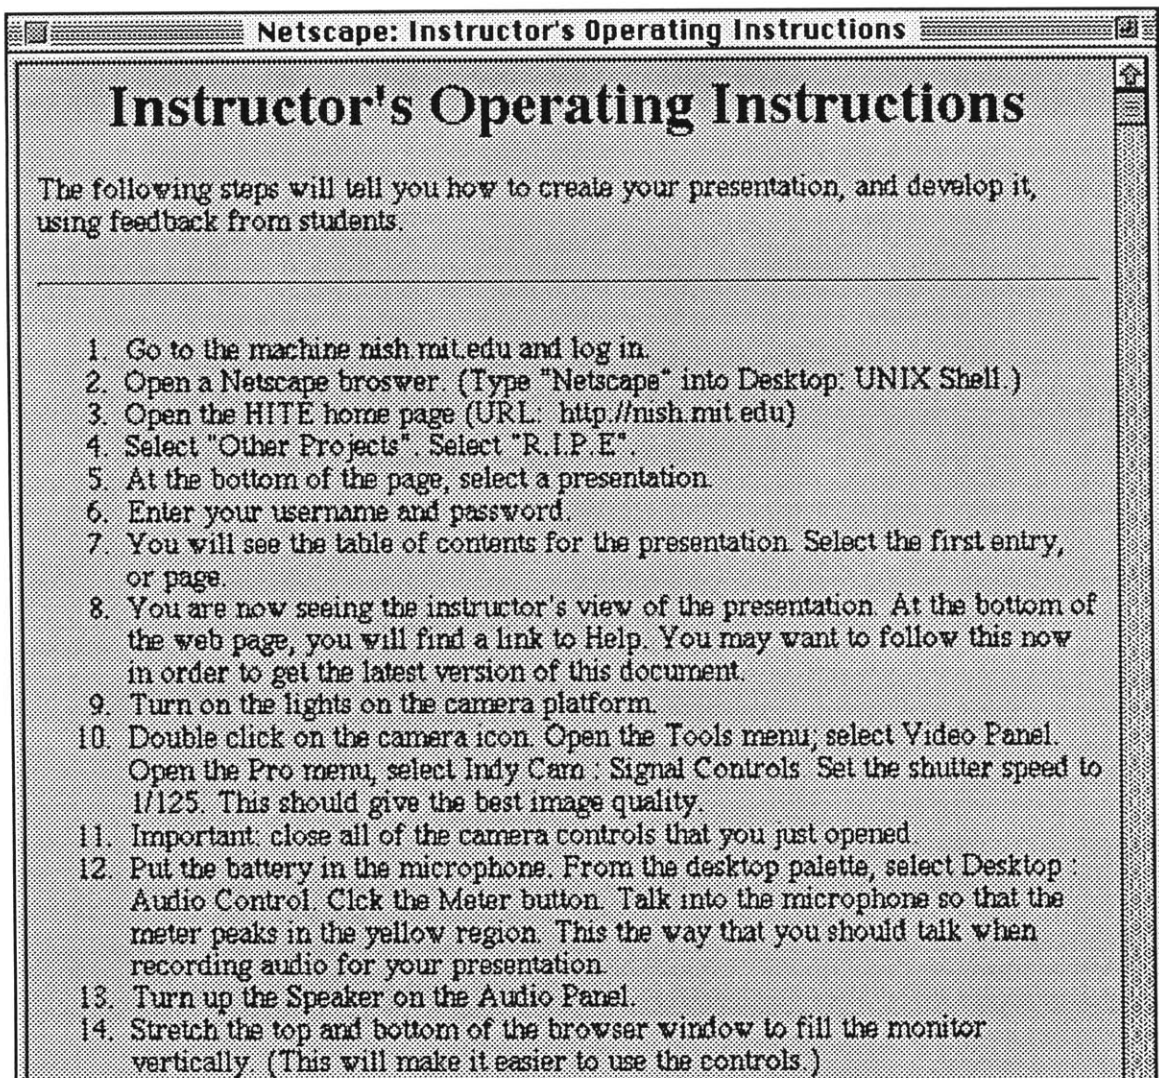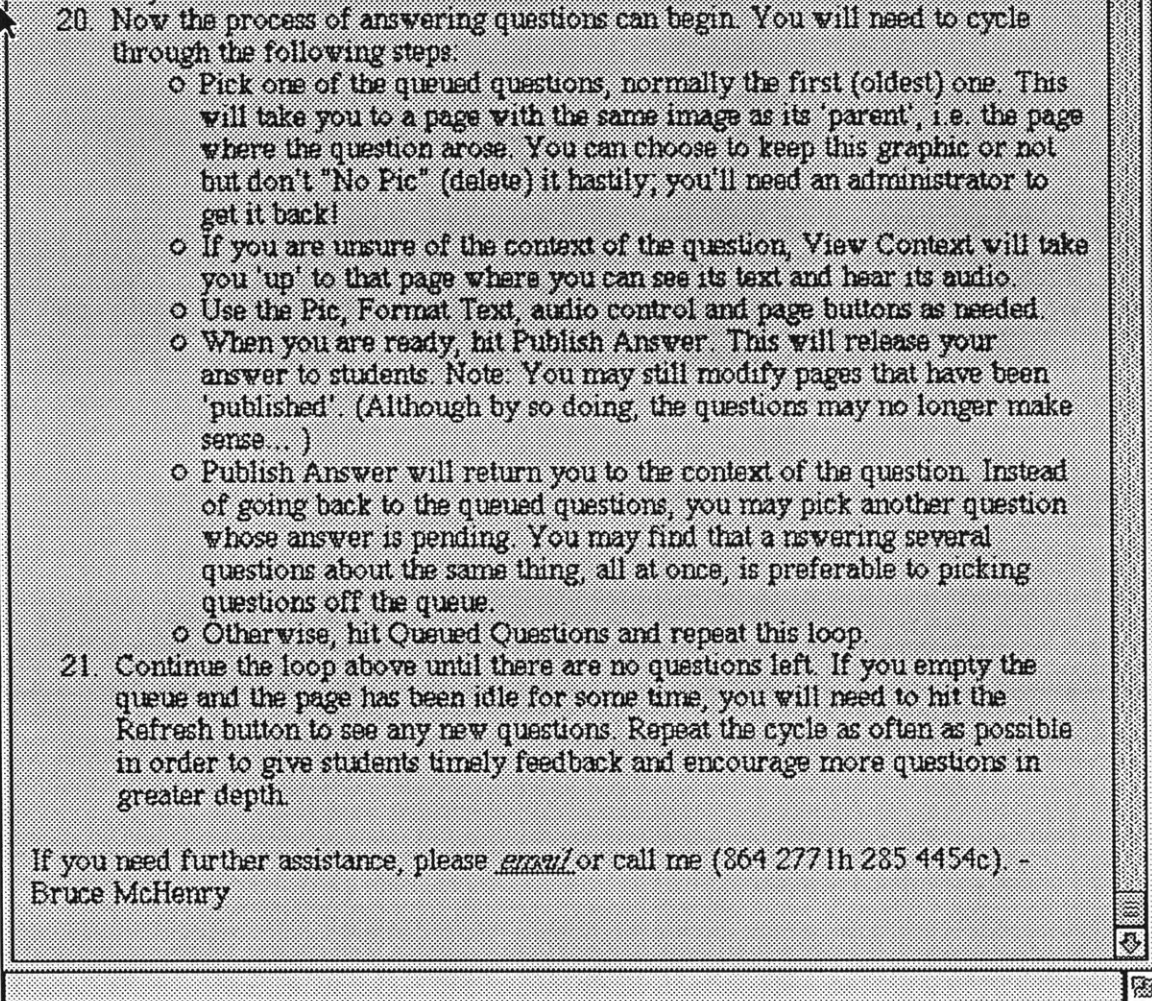
Figure 15: Instructors' Help (top)

15. You are ready to start authoring. If you are just starting to create a fresh presentation, the page will be empty except for the controls:
    o Picture will capture an image from the camera.
    o No Pic will delete the image being displayed.
    o Format Text will allow you to preview text typed into the large pink box above. The formatter will respond to HTML control sequences so, for instance, if you wanted to separate paragraphs, you would type <br> or

. The formatted text appears at the bottom of the page.
    o Record, Stop and Play apply to audio and have the obvious connotations. There is no easy way to edit your recordings so, if you make a mistake, simply Stop and Record again. After you hit record, wait a half second before beginning to speak.

16. Use these controls to create the first page of your presentation using any combination of speech, image and text. You can create multiple pages with the Next Page button. However, do not use this button unless you intend to add a page to the presentation. (Extraneous pages can be deleted but this is currently an administrator task. Also, there is no easy way to insert pages in the middle of sequence.)

17. When you are done, hit "Publish". This will release the pages for students to see (but you can still change them).

18. Tell the students that the pages have been published and can be accessed through the HITE home page where they can select [other projects]: R.I.P.E.

19. When students view your presentation, their questions will accumulate in a queue. The Queued Questions button (one of the instructor controls) will take you to this list of students' questions, which are listed in the order in which they were received.

Figure 16: Instructors' Help (middle)

20. Now the process of answering questions can begin. You will need to cycle through the following steps.
    o Pick one of the queued questions, normally the first (oldest) one. This will take you to a page with the same image as its 'parent', i.e. the page where the question arose. You can choose to keep this graphic or not but don't "No Pic" (delete) it hastily; you'll need an administrator to get it back!
    o If you are unsure of the context of the question, View Context will take you 'up' to that page where you can see its text and hear its audio.
    o Use the Pic, Format Text, audio control and page buttons as needed.
    o When you are ready, hit Publish Answer. This will release your answer to students. Note: You may still modify pages that have been 'published'. (Although by so doing, the questions may no longer make sense...)
    o Publish Answer will return you to the context of the question. Instead of going back to the queued questions, you may pick another question whose answer is pending. You may find that answering several questions about the same thing, all at once, is preferable to picking questions off the queue.
    o Otherwise, hit Queued Questions and repeat this loop.

21. Continue the loop above until there are no questions left. If you empty the queue and the page has been idle for some time, you will need to hit the Refresh button to see any new questions. Repeat the cycle as often as possible in order to give students timely feedback and encourage more questions in greater depth.

If you need further assistance, please _email_ or call me (864 277 1h 285 4454c). - Bruce McHenry

Figure 17: Instructors' Help (bottom)