To appear in Proceedings of the ACM Special Interest Group on Computer Science Education conference on Information Technology in Computer Science Education, 2006

Scaling Up a Distance Education Program In Computer Science

J. Mark Pullen
George Mason University
Dept. of Computer Science MS4A5
Fairfax, Virginia 22030, USA
+1.703.993.1538
mpullen@gmu.edu

ABSTRACT
This paper describes solutions to issues in presenting a graduate program in distance education online. The author acquired significant experience in synchronous Internet distributed education by experimental teaching online. This process has produced an open-source software system for synchronous online teaching and learning, Network EducationWare. The paper describes that system and the experience gained in expanding its use from one course to a graduate Computer Science degree program. Detailed discussions elaborate on problems encountered and their solutions. Experience has shown that online teaching in this format is beneficial to students, supportable by faculty, and affordable by the educational institution, using the effective implementation described.

Categories and Subject Descriptors
K.3.1 **[Computing Milieux]:** Computer Uses in Education – distance learning

General Terms
Design, Economics, Reliability, Experimentation, Human Factors

Keywords
Internet distance education, course management, education support

1. INTRODUCTION
For the past decade, the author has been engaged in expanding accessibility of courses in Computer Science at George Mason University (GMU) by synchronous Internet delivery [5,6]. The factors that are leading to considerable growth in distributed education worldwide [11] are strongly present in our Northern Virginia region. In the last five years, these led to an increasing number of faculty colleagues at GMU joining the online teaching effort, which recently reached a level of maturity that enables a student to complete a full Master of Science in Computer Science (MSCS) degree online [10]. During this process, the organization and delivery technology grew from a personal teaching practice to an operation supporting twenty courses, taught by sixteen faculty members. This in turn required development of a methodology for teaching students online and in the classroom simultaneously, which we call *simulteaching*, and a significant effort to create a support environment that can sustain a range of teaching styles at low cost, which is shown in [2] to be the lowest-cost approach to Internet distributed education. This paper reports on what we learned during the process of scaling up and describes the multimedia software and Web pages that we developed and have made available on an open-source basis. The paper begins with a description of the challenge in scaling up; this is followed by descriptions of the Web-based delivery technology we have developed and of the solutions we have found to the scaling problems.

2. ISSUES IN SCALING FROM AN EXPERIMENT TO A DEGREE PROGRAM
In the early 1990s, the author was responsible for a research project using advanced networks for distributed military training. At that time the Internet was, in most cases, not capable of supporting synchronous interaction. Those early experiments led to a conviction that nearly all the benefits of any social learning environment, from that of one-on-one-tutoring to that of large lecture halls, can be provided to remote participants using modern networking technology. As the capability of the Internet grew, it was appealing to test this vision in courses at GMU. Beginning in 1995, my students had the option of participating in class remotely. Initially we used Internet multicasting tools as described in [3]; later we experimented with an early commercial tool for delivery [7]. More recently we have integrated open-source multimedia tools into the system described in Section 3 below. This system, Network EducationWare (NEW) has proved quite successful. Students are very appreciative of the time saved by not commuting to class through our region's legendary traffic congestion. Also many faculty members are motivated to teach online in order to provide better accessibility for students, if the time demands to do so are not large. [13] reports that the synchronous mode is growing faster than any other approach in the dynamic field of distributed education. These positive responses encouraged us to support more courses; however, that required development of a much more mature support environment.

Growing from an online teaching capabiity that was essentially a hobby for the instructor to a capability that can reasonably support tens of courses and hundreds of students has proved to be a significant challenge. It has been necessary to take an operational perspective on all aspects of the delivery environment, integrating a production-quality database and tailoring a Web-based course management system. The work was shaped by several influencing factors:

- There was at least ten times as much of everything to be dealt with. The number of students, instructors, supporting materials and recorded sessions quickly became overwhelming. Our typical course will have up to 30 students in the classroom and 15 more online; currently we are supporting a total of 338 classroom students and 109 online students in ten courses.
- Available funding support was minimal. It was necessary to seek any source of support, including related research activities and administrative support staff. As the program grew, we were able to obtain limited funding from GMU based on cost savings in required classroom space.
- Faculty members are busy people; they generally do not have time available to take on extra roles to support distance education delivery. We were able to obtain some limited course release for faculty preparing first-time course offerings, but it also has been necessary to create a support system that places minimal extra demands on the faculty and supports their work with automation.
- Staff to support the activity necessarily is limited. Initially, the program was supported out of internal resources, building on software created in government-sponsored research work. While the software has continued to benefit from research spinoffs, operation has been transferred to the staff of a teaching lab operated by our college (two graduate students plus a few hours per week of a professional IT manager). We have automated as much of the support process as possible through a Web portal used by student, faculty, and supporters alike.
- Production users are much less forgiving of the technology than experimental users. The system must be made very robust, with fallback for any conceivable failure. We continue to discover possible failure modes and create ways to work around them.

3. NETWORK EDUCATIONWARE SYSTEM AND COMPONENTS

NEW is the third generation of teaching software we have used. Learning from the first two, we set out to create an open-source suite of software that relies as much as possible on quality software tools created by other parties [12]. Experience had taught us that the following properties are essential for simulteaching:

- Quality support for at least audiographics (voice, slides and real time annotations), using open media standards; video also is desirable, although not essential [8].
- Ability to support audiographic streaming over standard 56 kb/s modems connections to the Internet.
- Effective support to teach students simultaneously in the classroom and online.
- Audiographics recorded during class and accessible from a server, either by streaming or file download, using the same client suite as live access.
- Web-based access and software load/checkout for ease of operation.
- Availability on multiple platforms for ubiquitous use (Windows, Linux and Macintosh systems).

The NEW system software that we have assembled to meet these requirements is represented in Figure 1 and summarized in the following sections. Executable and source code for all system components is openly available for academic purposes. (By way of contrast: A commercial system that can support similar requirements, Microsoft Live Meeting, currently lists annual license price of $3,000 per server plus $180 per named user [4].)

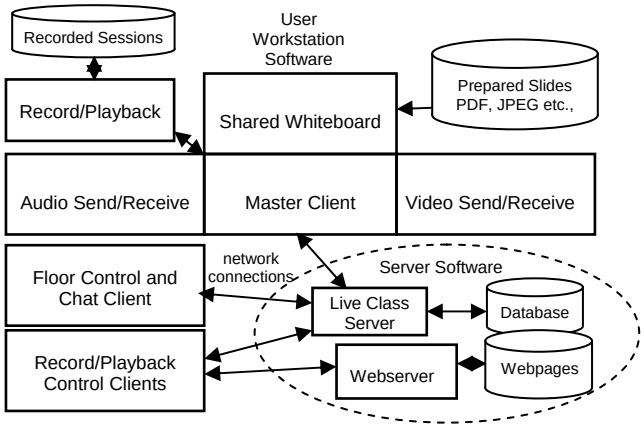

Figure 1. NEW system software architecture

3.1 Client Software

The multimedia interface software in NEW derives from a variety of sources and runs on Windows and Linux-x86 platforms, with a Macintosh capability planned for late 2006. The composite user interface for all tools is shown in Figure 2.

We consider the *Speak Freely Internet Audio* software to be the most important component in the NEW system, both because it is essential to the students' learning experience and also because conveying voice with good quality over the Internet at low data rates presents a big challenge. SF is capable of passing good voice quality over the Internet, using a standard sound interface, and requiring only 20 kilobits per second of network capacity. We have added a graphic interface that provides all needed user functions in one easy-to-use package.

The *Whiteboard* provides the other key element for teaching online: graphics. It will display a precomposed graphic prepared in several open formats: text, HTML, JPEG, and Adobe Portable Document Format (PDF); and it will convert LaTeX, OpenOffice, and Microsoft PowerPoint formats to PDF and JPEG automatically. The precomposed graphics can be annotated during class with lines, rectangles, ellipses, handwriting, and text in any color, a very useful feature for maintaining the attention of the visual learner. We prefer to use the whiteboard with a Tablet PC interface so that it becomes a surrogate chalkboard.

The optional *Video* tool is capable of multiple network formats, including standard H.323 conferencing. A typical delivery rate for

NEW is two frames of 320 by 240 pixels per second, although rates up to 30 frames per second are possible. While we have found that, for teaching Information Technology, video provides a marginal benefit at relatively high cost, as reported by Pullen (2001), we offer it as an option to students who have high quality Internet service.

The *Master Client* encapsulates data from the multicast applications into TCP tunnels to the Live Class server, prioritized according to the importance of each multimedia tool (audio first, whiteboard second, video last). It can support a viable class connection over a 56 kb/s modem, without video. The combination of clients and their network configuration established by the master client is controlled by a configuration file downloaded from the supporting webserver at the beginning of a NEW session. If software updates are indicated, the master client also downloads and installs them.

The NEW *Floor Control* shows the participants in the session, controls access to the virtual classroom "floor," provides for text questions to the instructor and text chat among the participants, and accepts URLs from the floor holder for browser launch on all participating client systems. It supports a "virtual hand raising" mode for lectures and an "anyone can have the floor" mode for seminars and meetings.

NEW *Record and Playback Clients* control their respective servers. They feature VCR-like button icons and an elapsed time readout. The playback control also is capable of jumping forward and backward to the next slide in the presentation. Recordings require about 5 Megabytes of disk per hour of class.

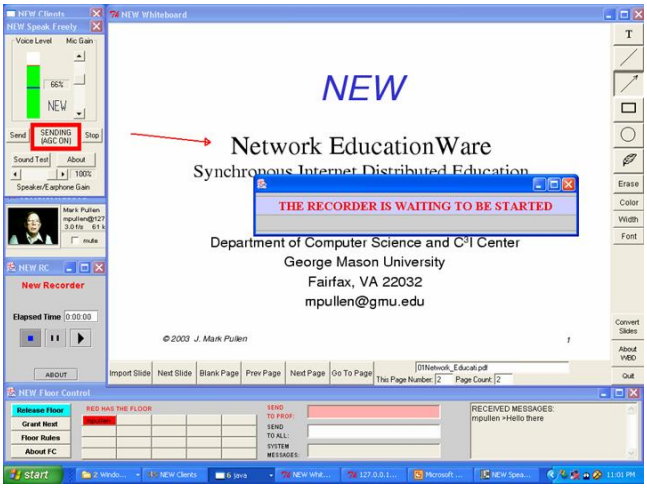

Figure 2. NEW user interface for open-source synchronous teaching

3.2 Server Software

The *Live Class Server* is the core of the NEW system. It implements group communication over the general Internet among a group of participating workstations by accepting a data stream combining transmissions from multicast conferencing tools on the floorholder's workstation and sending copies to all participants' workstations. It provides access control using either the system database or an external authentication service. As delivered from our website it will support 20 users on a 1 GHz Linux system; however, it can be configured for up to 50 users on a 3 GHz system.

The *Record* and *Playback* servers are used to create and play streaming recordings that capture the information sent by NEW clients over the Internet from the instructor's workstation. Playback is accessed through the same software suite and also can be performed offline by downloading the recording files. Each segment of an online playback can be Web-linked as a URL.

The NEW Course Management Webpages are the key to effective management of the mass of detail involved in supporting multiple courses. They provide teaching and learning functions, support and course management functions, authentication, and system administration functions. A single portal page provides access to all of these facilities, as described in [9]. Our webpages are hosted on an Apache Webserver that supports the PHP language needed for our webpages and MySQL or any database system supporting the standard Structured Query Language (SQL). Web-based support provides ubiquity and portability. It also makes possible data access over the network that we use to implement the chat room feature. The instructor version of our portal's "Welcome" page is shown in Figure 3.

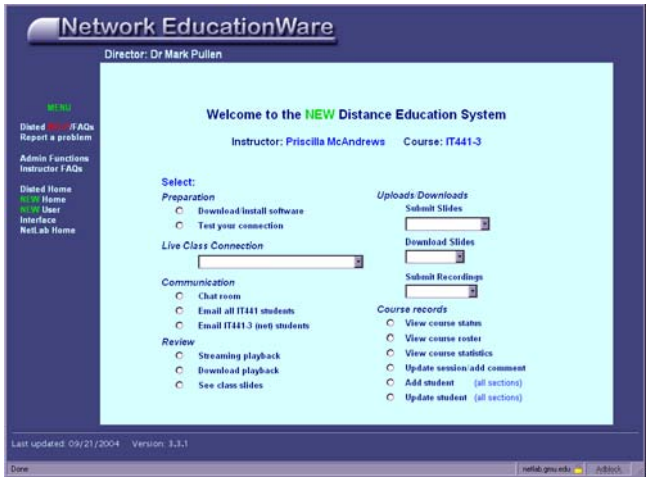

Figure 3. NEW Web Portal "Welcome" page

4. SCALING ISSUES AND SOLUTIONS

In this section we describe the practical solutions we have developed for scaling up.

4.1 Teaching Platform

Although it also can be used with a mouse, most instructors find the NEW whiteboard is most easily used with tablet input. After trying several configurations, we have found that the most convenient way to provide this is a "Tablet PC" laptop computer, which can provide the complete audiographic interface in one robust unit and also costs less than a computer with separate tablet. We optionally add to the Tablet PC an external "Webcam" video camera, full size keyboard, and wireless microphone. The whole package costs around $2000. It is kept in the supporting

teaching laboratory to be delivered and plugged in at class time by a student supporter. While in the laboratory, it is safe from poking fingers and also is available to perform software maintenance and NEW file upload/download. We keep a spare Tablet PC fully configured as a backup.

4.2 Enrollment Management

At last count our student database had over 2600 entries for individual course enrollments over a three year period. (This includes classroom students, who also access NEW for recordings.) While the faculty try to deal with each student as an individual, this is simply not possible for the small support staff. To keep track of the numbers, the NEW Course Management system maintains a database with entries for every student in every course and every login for any purpose. This allows automatic production of reports (for instructors and staff only) showing participation, both individual and as overall course statistics. We get the course enrollments from the registrar at the beginning of each semester; thereafter, the instructors maintain the enrollment records, update passwords, etc. because they are the ones who accurately know each student's status. In the past we have generated our own passwords and sent them to the students via campus email; now we also have the capability to use either the campus Lightweight Directory Access Protocol (LDAP) server or the campus Post Office Protocol (POP) email server to verify against centrally administered passwords.

4.3 Assistance for Students

It is important to have a way for all users to get assistance. We provide an email address where help can be obtained, and callback for issues that require discussion. Happily, most students are able to load the NEW clients without assistance. The client may be downloaded through the NEW portal webpage and connected to a test server through the same page. This avoids many problems, and allows students to verify that their client will work when class time arrives. However, there always are a few users who have problems. The first line of assistance here is the supporting laboratory staff, who operate their own clients to monitor outgoing class availability and thus are familiar with the software and supporting webpages. In particularly problematic cases, the user is referred to a system administrator who can resolve almost any problem. On rare occasions, a software bug is identified by this procedure, and returned to the development staff, who log it for future action and recommend a work-around procedure.

4.4 Assistance for Faculty

The same organization is available to assist the faculty before and during class. Before class, instructors upload slide files (in Adobe PDF, LaTeX, or Microsoft PowerPoint formats) to be loaded in The NEW whiteboard before class. The laboratory student staff load these files and also upload class recordings to the playback server. They also offer assistance in formatting slides for best readability. During class, the staff monitors the session to ensure that sound and slides are present and understandable. One support person can monitor two or three classes simultaneously. If there is a problem with the classroom equipment, this person is then available to help correct the problem. We run the whole program with two graduate students, who cover a total of 20 hours of class time per week in addition to providing faculty and student assistance.

4.5 File Management

The slides and recordings are kept on a Linux server in the secure laboratory server room. As described above, they are uploaded and downloaded via the webserver. The NEW Web Portal provides status of all files, by class session, to instructors and staff via Web browser. Backups are made regularly from the fileserver (including the NEW database) by the laboratory system administrator. This is all very mundane, but it also is essential for effective system operation.

4.6 Scheduling

The distance education program imposes three new constraints on the course and classroom scheduling process: (1) Courses using NEW must be in classrooms with computer projectors and good network connections; (2) The number of simultaneous courses using NEW cannot exceed the number supportable by available laboratory staff, servers and Tablet PCs; (3) Online courses in the MSCS program are scheduled in an announced rotation so that a student can plan a degree program in advance, subject to cancellations only for emergencies.

4.7 Asynchronous Online Course Integration

In principle, assuming a complete set of online supporting materials, a student could complete a course from the recordings made in the NEW system. Our courses all have webpages that provide access to supporting materials; some also use a Learning Management System (LMS). GMU supports WebCT; Blackboard and the non-commercial Moodle are other LMS examples. Currently, we are experimenting with using this asynchronous support to offer courses outside the normal schedule. Initial student reaction has been enthusiastic, although students have little interest in working on asynchronous courses during normal semesters. They want the flexibility of asynchronous courses during the Summer Term, when very few graduate courses are available. The author has offered pilot courses for the last two years experimentally; that experiment was sufficiently successful that two other faculty members will join in testing the concept in summer 2006. The instructor, who originally presented the recorded course, is available via email and chat room to mentor and provide assignments and examinations for the asynchronous student. The department faculty has agreed that, under these conditions, asynchronous offerings are acceptable from an academic standpoint. The other issues we had to resolve were administrative and came about because the mode of delivery does not fit the normal Summer Term enrollment process: publicizing the courses, paying the instructors on a per-student basis, and arranging for individual payment.

4.8 Meetings Outside Class Hours

As at most institutions, GMU instructors have office hours when they are available for interaction with students. We have set up chat rooms, accessible through a Web browser interface, where students from a particular course can meet at any time. We support this with the open source tool phpMyChat; similar capabilities are available in any LMS.

4.9 Administrative Support

Carswell [1] argues persuasively that is important that some arrangement be available to provide the full range of academic services to distributed education students. As described above, it is necessary for support management and scheduling to cooperate so that the student can access the course with no impediments. It also is necessary to consider such issues as registration, bookstore access, library access, and proctoring examinations for online students. We have decided that, absent special arrangements between student and instructor, we will provide support for regional online students. By this we mean students who can and do come to campus for examinations. Most of our students live within a two hour drive of our campus; however, a few do make special arrangements. (At the time of writing, one student is participating from Belgium and another from Lebanon.)

4.11 Software Enhancement

As the developer of NEW, the author has another problem: how to ensure that the software is acceptably free of bugs when it is deployed. The solution to this problem is a traditional one: pilot implementation with a test group of users consisting of the author's own students. Their aid is enlisted to find and fix bugs; as Computer Science students, many of them participate enthusiastically.

5. CONCLUSIONS

We have found that offering our MSCS courses synchronously over the Internet is well received by students because it increases accessibility greatly and thus adds usable hours to the student's day. The process is not difficult for students and faculty, but it does require support to scale up to a full degree program. Due to the NEW Web Portal and a thorough suite of supporting automation, this support is affordable. The key to success is to be aware of what needs to be done and deal with these needs proactively. This paper has summarized the procedures we have used to achieve success in that process. A few institutions beyond GMU have begun to experiment with NEW. We are optimistic that the next level of scaling will include its much broader adoption. The software and supporting documentation are available at http://netlab.gmu.edu/NEW.

6. ACKNOWLEDGEMENTS

NEW owes its existence to the dedicated labors of many students, and particularly to that of the GMU C^4I Center Networking and Simulation Laboratory staff. Priscilla McAndrews, Charles Snow, and Nicholas Clark deserve special mention. Many others are included on the NEW website listed above.

7. REFERENCES

[1] Carswell, L., The "Virtual University": Toward an Internet Paradigm?, *ACM SIGCSE Bulletin 30*, 3, (Sep. 1998), 46-50

[2] Goodwin, C. and Bowman, M., Is the bottom line of online out of line? Calculating the total cost of online courses in Technology Curricula, *Proceedings ASEE Annual Conference 2004* (Salt Lake City, UT, June 2004)

[3] Macedonia, M. and Brutzman, D., Mbone Provides Audio and Video Across the Internet, *IEEE Computer 27*, 4 (Apr. 1994), 30-36

[4] Microsoft Corporation, Office Microsoft Live Meeting Retail Pricing, http://www.microsoft.com/office/livemeeting/howtobuy/default.mspx

[5] Pullen, J., Synchronous Distance Education and the Internet, *Proceedings Internet Society Annual Conference 1998* (Geneva, Switzerland, July 1998), published online at http://www.isoc.org/inet98/ proceedings/ 4b/4b_1.htm

[6] Pullen, J., The Internet Lecture: Converging Teaching and Technology, *ACM Special Interest Group on Computer Science Education (SIGCSE) Bulletin* Vol 32 No 3 (Sep. 2000), 101-104

[7] Pullen, J. and Benson, M. (1999), ClassWise: Synchronous Internet Desktop Education, *IEEE Transactions on Education 42*, 4 (Nov. 1999); special multimedia CDROM issue, printed abstract 370

[8] Pullen, J., Applicability Of Internet Video In Distance Education For Engineering, *Proceedings IEEE Frontiers in Education 2001* (Reno, NV, October 2001), T2F-14–T2F-19, online at http://fie.engrng.pitt.edu/ fie2001/papers/1242.pdf

[9] Pullen, J. and McAndrews, P., A Web Portal for Open-Source Synchronous Distance Education, *IASTED Journal on Advanced Technology for Learning 2*,1 (Jan 2005), International Association of Science and technology for Development, Calgary, AB

[10] Pullen, J., Simon, R. and McAndrews, P., An Online Graduate Computer Science Program Delivered Via Simulteaching, *Advanced Technology for Learning* Vol 2 No 3 (Jul. 2005) 148-155, ACTA Press, Calgary, AB

[11] Sloan Consortium, Sizing the Opportunity: The Quality and Extent of Online Education in the United States, 2002 and 2003, online at http://www.sloan-c.org/resources/overview.asp

[12] Snow, C., Pullen, J. and McAndrews, P. (2005), An Open-Source Web-Based System for Synchronous Distance Education, *IEEE Transactions on Education 48*, 4 (Nov. 2005), 705-712

[13] Wilson, J., After the Fall: Lessons of an Indulgent Era, plenary presentation, Distance Education 2003, University of Wisconsin, unpublished, available online at http://www.jackmwilson.com/ArticlesTalks/eLearning-Wisconsin2003.ppt