Situating CoWeb: a Scholarship of Application

Jochen Rick and Mark Guzdial
College of Computing / GVU Center
Georgia Institute of Technology
{jochen.rick, guzdial}@cc.gatech.edu

Abstract

Since 1998, we have been developing and researching CoWeb, a version of Ward Cunningham's WikiWikiWeb designed to support learning. In this article, we summarize our results of situating CoWeb across the academic landscape of Georgia Tech. In architecture, CoWeb enabled faculty to serve more students in a design-based course. In English composition, a comparison study demonstrated significant learning benefits without incurring disproportionate costs. Yet, situating CoWeb was not always successful. In many STEM (Science, Technology, Engineering, and Mathematics) classes, students actively resisted collaboration. From these studies, we conclude that the culture of the classroom and the discipline needs to be compatible with the medium for CSCL (Computer-Supported Collaborative Learning) to be effective. Finally, we demonstrate how collaboration can be designed into the culture. A new class on introductory computing was explicitly designed to take advantage of the collaborative possibilities that CoWeb affords. We report our findings of the success of this approach.

We situate this research as a scholarship of application. We demonstrate that this mode of scholarship is an important mode of scholarship in the learning sciences. Unlike traditional scholarship of discovery, we are not solely concerned with discovering new knowledge. Instead, we support others in the application of a new technology to serve genuine and complex learning situations. By doing so, we seek to understand the potential that one new medium (WikiWikiWeb) has for supporting learning.

Understanding a New Medium

Media, the channels through which we forge and communicate meaning, affect us as individuals and as a society (McLuhan, 1964; Meyrowitz, 1985). New media change how we relate to ideas, to others, and to ourselves (Turkle, 1984; Turkle, 1995). As such, a new medium has the potential to be a powerful and natural learning environment (Rick and Lamberty, 2005). Unfortunately, realizing that potential is challenging (Bolter and Grusin, 1999). The effect that a new medium has on society and individuals is frequently not envisioned by its creator; typically, it is only understood in historical hindsight (McLuhan, 1964).

When Gutenberg started printing Bibles, he had no idea that the medium would eventually lead to a loss of power for the Medieval Church (McLuhan, 1962). As Bibles became more accessible, the church no longer had a monopoly over religious information (and thereby salvation). People could read the Bibles in their own native language by themselves. The printing press enabled Luther, and the Church was changed forever.

When Edison invented the phonograph, he envisioned a "talking machine" that could be used for correspondence and dictation (Radick, 2003). Edison's original phonograph could both record and playback audio. Market forces, however, wanted something different. Commercially, the phonograph only succeeded as a playback machine.

The desktop user interface (i.e., overlapping windows, mouse, menus, etc.) was invented at Xerox's Palo Alto Research Center (PARC) in the 1970s. Yet, the personal computer explosion driven by that interface in the 1980s was not driven by Xerox. Why? When Xerox executives were shown the new personal computer, they failed to recognize its commercial potential (Stone, 1998). Some feared that a personal computer spelled doom for paper, Xerox's key business. Xerox failed to capitalize on the opportunity and others introduced the world to the personal computer.

As these examples demonstrate, understanding a new medium is a difficult undertaking. It often requires significant use and evolution of the medium itself. In this article, we summarize our efforts to understand the learning potential of one new medium, Ward Cunningham's WikiWikiWeb (or Wiki). In November 1997, we started developing our own wiki implementation, CoWeb (Collaborative Websites), to serve learning purposes. Since then, we have worked with faculty in other departments to deploy CoWeb to serve learning needs.

In this article, we situate CoWeb in two ways. First, we situate CoWeb (our research) theoretically. Extending traditional notions of scholarship, we situate our work as what Boyer (1990) terms a *scholarship of application* in design-based research. We apply the new medium to practical learning situations. Thereby, we hope to both support learning in these situations and to better understand the medium. Second, we situate CoWeb into different learning situations. We describe four different attempts to situate CoWeb in learning situations. We reflect on our findings to gain a better understanding of both our medium and learning in situ.

Categorizing Design-Based Research

Recently, design-based research (DBR) has received considerable attention from the learning sciences, earning special consideration in both *Educational Researcher* (vol. 32, no. 1) and *The Journal of the Learning Sciences* (vol. 13, no. 1). DBR is a paradigm for the study of learning in context through systematic design and study of that design in practice (Collins, 1992). Two parts are essential—iterative development and authentic contexts. In an iterative development cycle, the design informs the research and the research informs the design. Acknowledging that crucial aspects of learning are missing in a laboratory context (Lave, 1997), DBR is situated in authentic contexts, such as the classroom (Barab and Squire, 2004). DBR aims to make sense of the complexity inherent to the authentic context.

DBR has its roots in design experiments (Brown, 1992). Often, the terms design-based research and design experiments are used interchangeably; unfortunately, this undermines much of Brown's contribution (McCandliss et al., 2003). Brown (1992) narrowly defines design experiments as iterating between the laboratory and the classroom. Work in the classroom informs the researcher on what to study in the laboratory. Work in the laboratory informs the design of the classroom. In that formulation, design experiments can be categorized as design-based research. But, not all design-based research can be categorized as design experiments.

Since its inception (Collins, 1992; Brown, 1992), DBR has grown into a broad classification, encompassing a wide variety of research paradigms. Design experiments is just one. In this paper, we introduce another—situating a new medium. We categorize this paradigm for DBR as a *scholarship of application*.

Scholarship in the Learning Sciences

Traditionally, scholarship has been defined as original research that focuses on creating new knowledge for its own sake (Boyer, 1990). Boyer terms this mode the scholarship of *discovery*. To this traditional mode, he adds three additional modes of scholarship that are relevant to the modern academy: scholarship of integration, scholarship of application, and scholarship of teaching. Scholarship of *integration* is concerned with integrating knowledge by forging connections across disciplines. Instead of creating original research, the focus is on bringing new insight to bear on original research. Scholarship of *application* is concerned with applying knowledge to consequential problems and reflecting on that endeavor. Scholarship of *teaching* is concerned with teaching the knowledge of your field to others.

While most learning sciences research conforms to the traditional discovery mode, the other modes are already represented. First, learning sciences is located at the intersection of several disciplines—education, psychology, computer science, etc. As such, there is a need for scholarship that integrates the disciplines. Next, as learning research often seeks to teach the skills and practices of another discipline, the research can be viewed as scholarship of teaching from that discipline's perspective. So for instance, research on math learning may be considered a scholarship of teaching from mathematics's perspective. Finally, education is a consequential problem. A successful application of learning theories to solve an important learning problem is a significant contribution even if the solution is too situated for new knowledge be abstracted.

It is in this mode, that we categorize our research. As in other design-based research, we start with an initial design. Then, we apply our design to solve consequential learning problems, iterating simultaneously on both the design and the authentic context of use. By reflecting on the results, we complete a scholarship of application. While this description might fit a scholarship of discovery, our contribution is different. Instead of trying to discover important abstract knowledge, we seek to create and understand an important concrete design.

DBR: Discovery vs. Application

Design-based research falls on a continuum from being a scholarship of discovery to being a scholarship of application. As in other disciplines (Boyer, 1990), the learning sciences has traditionally valued discovery more. For instance, design experiments are in this mode: The goal of a design experiment is to abstract new knowledge that can be applied elsewhere (Brown, 1992). Yet, learning is a situated enterprise (Lave, 1997; Lave and Wenger, 1991).

There are significant, complex learning problems that are not served by the discovery mode. These problems and their solutions are too complex and situated to allow researchers to abstract from the context. Scholarship of application research aims to solve these problems and (to make it research) reflect on the solution.

In the discovery mode, whether the design succeeds is less important than whether new knowledge can be discovered in the context. Barab and Squire (2004) recall a design experiment that showed significant improvements in learning, but the course was canceled because it was too costly. While the design experiment succeeded (in discovery), the course failed (in application). In contrast, in the application mode, the success of the design is primary.

Scholarship of application DBR acknowledges that practical considerations too are important. For a design to be sustained, it needs to be economically viable—the benefit must be worth the cost. It is necessary to consider issues of usability, scalability, and sustainability (Fishman et al., 2004). Confining time and effort is crucial to a design's success. For instance, using adult volunteers in CSCL has received considerable attention; however, as was made clear by the ICLS 2000 panel "Scaling Educational On-line Communities: The Role of Volunteerism in Doing Large-Scale Educational Projects On-line," the costs to recruiting and organizing volunteers is often substantial and prohibitive. Likewise, scaling a solution is not trivial (Songer, 1998).

As it tries to solve significant, complex problems, research in the application mode tends to be larger in scope than research in the discovery mode. Because of its smaller scope, the discovery mode retains more control over its context. The context can be controlled in such a way as to abstract new knowledge. In contrast, results of a scholarship of application are more situated. The situated solution (i.e., the concrete design) is the contribution. In the discovery mode, the contribution is determined by the applicability of the new knowledge. In the application mode, the contribution is measured by the applicability of the concrete design.

In this article, we detail one instance of a scholarship of application. We demonstrate the value of our solution and its applicability to learning. Through our solution (and our reflection), we make a significant contribution to the learning sciences. Thereby, we demonstrate that the application mode can be useful for design-based research in the learning sciences. While we value scholarship of discovery, we feel it is important to acknowledge scholarship of application.

Situating a New Medium

New media change how we relate to ideas, to others, and to ourselves. As such, a new medium has the potential to be a powerful and natural learning environment. The computer is a particularly flexible tool to create new media (Kay and Goldberg, 1977). Consequently, many (Papert, 1993; Resnick et al., 1996; diSessa, 2000) have sought to harness this potential (of creating new media) for learning. In the field of CSCL, particular learning potential is found in computer networks, such as the Internet (Koschmann, 1996; Scardamalia and Bereiter, 1991). Through the network, learners can connect with information, peers, and experts. We, the authors, are computer scientists. As such, our contribution to the learning sciences tends to focus on the potential of the computer and the network.

In this article, we concentrate on understanding the potential of one new medium, the WikiWikiWeb. As detailed in the next section, this medium has great potential to support collaborative learning. As with other new media, realizing that potential is far from trivial. Significant use is necessary to properly understand the potential of the new medium (Bolter and Grusin, 1999; McLuhan, 1964). Adopters need the freedom to innovate. The more contexts and uses the medium is applied to, the more likely it is that the true affordances of the new medium will emerge.

As such, this research cannot be conducted within the limited scope and tight control of a scholarship of discovery. Instead, we opt for a scholarship of application. Our research aims are not to abstract new knowledge, but to better understand the medium. What are its affordances? In which situations and uses does it succeed? In which situations and uses does it fail? To answer these questions, the large scope and loose control of a scholarship of application are actually to our advantage. They help us better understand the new medium. Just as design experiments are a useful category of DBR in the discovery mode, we see *situating a new medium* as a useful category of DBR in the application mode. The methods are quite different, but each makes important contributions to the learning sciences.

CoWeb: a CSCL Wiki

The World Wide Web was designed to be a particularly flexible medium (Berners-Lee, 1999). The flexibility of hypertext combined with a world-wide network of computers makes it a natural platform for creating new media. Blogs (or weblogs) are one example (Mortensen and Walker, 2002). Wiki is another (Leuf and Cunningham, 2001). A wiki is a website that invites all users to edit any page and add new pages, using only a regular web browser. The text is edited in an HTML text area without special applets or plug-ins. Cunningham named the WikiWikiWeb after the Hawaiian expression meaning quick ("wiki wiki"). His reasoning was simple: The quickest way to create a web-site is to ask all visitors to be authors. The Wiki is an unusual collaborative space in its total freedom, ease of access and use, and lack of structure. As such, the Wiki has a significant potential to support collaborative learning.

In addition to having significant learning potential, the WikiWikiWeb has several practical advantages, making it worthy of a scholarship of application. First, the underlying infrastructure (access to the Internet and the Web) needed to support a wiki is becoming common. Secondly, as a purely server-side application, there is no need to distribute special software to users. All people need to access a wiki is an ordinary browser. This lowers the barrier to entry, making it easy for people to adopt. Consequently, the Wiki has become quite popular.[1] Third, the costs of setting up and maintaining a server are reasonable. A single mid-range server (around $2000 as of this writing) can support hundreds of classes and thousands of students. The most substantial cost is teacher time. Yet, CSCL has the potential to actually decrease that cost (Barab and Squire, 2004, p. 11). By engaging students in collaboration, the large numbers in classes can be leveraged to create greater opportunities for discussion, reflection, and (consequently) learning. Because the increased opportunity

[1] According to Alexa (*http://www.alexa.com*), Wikipedia, a wiki encyclopedia, is one of the top one hundred sites on the Web.

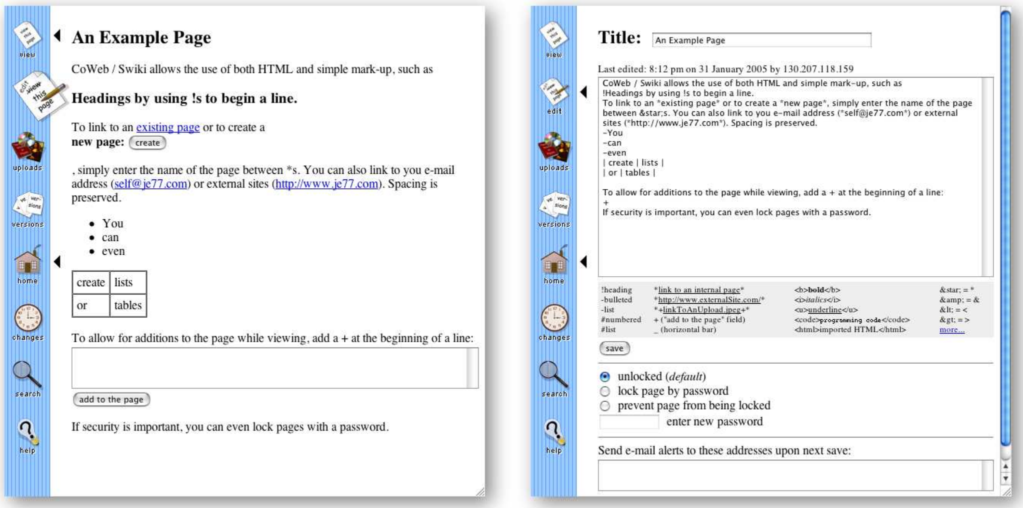

Figure 1: Viewing / Editing a CoWeb Page in a Standard Web Browser

for learning is coming from the students themselves, the cost for the teacher does not need to rise any further than simply providing oversight for the process.

Thus, for relatively low costs, we believed significant learning benefit could be gained. We designed CoWeb (our Wiki implementation) to understand and realize the learning potential of the WikiWikiWeb. Following an active DBR strategy, we evolved both the technology and its context of use simultaneously.

Like a wiki, a CoWeb looks like traditional static web-site, except that every page has a set of buttons that allow the user to do various things, such as edit the page (Figure 1), upload attachments, and view the history of the page over time. Links between pages are easily created by referencing pages within the same site by name (e.g., *Page Name*). If a page with the given name does not already exist, a “create” link shows up next to the name upon save; following the link creates the new page (Figure 1).

While maintaining the essential characteristics, CoWeb varies from the original Wiki implementation. CoWeb has had features added and the interface streamlined to fit well into classroom use (Guzdial et al., 2000). Changes were made to better serve our needs (i.e., supporting learning in our various settings). For instance, the editing format is a bit different. The original provides a simple text notation without HTML; CoWeb accommodates both the text notation and using HTML. In our academic community, most users know some HTML and want to use their prior knowledge. Also, CoWebs are easily adaptable. Many of our class CoWebs contain extra features that were designed for those specific sites. For instance, in mathematics and engineering, the text notation was augmented to allow users to easily share Matlab code. Other sites show “hot spots” at the top of every page; these can be update to draw attention to “hot” pages, such as the current discussion.

Situating CoWeb: Medium in Culture

In January, 1998, we introduced CoWeb to Georgia Institute of Technology (Georgia Tech). Since then, we have iterated on both our design and its context of use. We have seen adoption on a massive scale. Over 300 classes have used CoWeb as a collaborative space. This use spans the academic landscape of Georgia Tech, including Architecture, Chemical Engineering, Computer Science, LCC (Literature, Communication, and Culture), and Mathematics. At times, we have supported up to ten servers to support these classes. At the time of this writing, Chemical Engineering, LCC, Mathematics, and Bioengineering have and maintain their own CoWeb servers. Fueled by this adoption, the software has gone through 15 major iterations becoming more powerful, flexible, and robust. A wide variety of educational activities have been invented by teachers for their classes (Guzdial et al., 2001), and we have cataloged some 25 core activities that we see tailored to meet specific class needs (Collaborative Software Laboratory, 2000).

What follows are four case studies of CoWeb adoption. In each, we situate CoWeb in a different context. We describe how the use in context influenced our design and our understanding of the learning potential of the WikiWikiWeb. To ground each case study (Yin, 2003), we have included a rich description of each effort, detailing use, development, and findings. This rich description should allow readers to better understand the nature of our design-based research (Hoadley, 2002).

The case studies are presented in roughly chronological order. First, we situate CoWeb in architecture. The intense use in architecture drove much of the early understanding and evolution of CoWeb. Second, we situate CoWeb in English composition. We take a more traditional research approach, examining learning in a comparative study. We find that the instructors were able to adopt CoWeb successfully, improving learning without incurring disproportionate costs. Third, we situate CoWeb in STEM (Science, Technology, Engineering, and Mathematics) classes. In contrast to our success in architecture and English composition, results are disappointing. We conduct interviews and surveys to understand why. We find that there are deep cultural barriers in these classes that make adoption of CoWeb largely futile. Fourth, we situate CoWeb in Media Computation, an introductory computer science class organized around media manipulation. We face the same cultural barriers as in the other STEM classes. Yet, through an innovative design, we are able to overcome those barriers. We detail the essential role that CoWeb played in realizing this design.

A recurrent theme in these case studies is one of cultural compatibility—if the culture of the context is not compatible, the medium will not succeed. In architecture, both the culture of the field and the culture of the classroom are compatible with collaboration. CoWeb succeeds. In English composition, the culture of the field is not collaborative, but the culture of the classrooms we studied is. Again, CoWeb succeeds. In STEM classes, neither the culture of the field, nor the culture of the classroom are compatible with collaboration. Consequently, CoWeb fails. In introductory computer science, once again, the culture is problematic. Yet, we succeed by changing the culture. CoWeb played an essential role in allowing that to happen.

In Architecture: Design and Practice

One of the first adoptions of CoWeb was in the College of Architecture. In architecture, both the culture of the field and the culture of the classroom were a good match for a new collaborative medium. Architecture is a design field—the primary focus is on creating designs and sharing them with others. Architects must be able to convey their designs to fellow architects, clients, project managers, civil engineers, etc.

In the curriculum, it is understood that learning to design is an active process of doing (Craig and Zimring, 2000). Because of this, the architecture curriculum is seen as a model of how to teach design in an active manner (Kehoe, 2001). Classes are often organized as design studios: 6-15 students spend a considerable amount of time (e.g., 12 hours per week) working on different designs. Students frequently work in groups. In the studio, students work and informally discuss their work with each other and their instructor. At different stages of the design, they formally display their progress for critique by the instructor and fellow students; if available, external critics come in to critique the work. The feedback that students get from critique is fundamental to their learning (Schön, 1987). Schön (1987) characterizes design as a process of reflection-in-action. A designer must take an action and then reflect on that action; others are often part of both the action and the reflection step.

Communication is fundamental to the field and the classroom. So, when CoWeb came along, the architecture faculty were interested in applying it to serve their communication needs. Applying new communications technology in the architecture studio is by no means a new idea. What made CoWeb unique was that it was a fairly unsophisticated (low bandwidth, no special hardware) technology (Zimring et al., 2001); other efforts focused on high-resolution video and new hardware that make widespread adoption of the technology impractical.

Working together with us, faculty created the CoOL Studio (Collaborative On-Line Studio for Architecture) CoWeb for a graduate-level studio focusing on courthouse design (Craig and Zimring, 2000; Zimring et al., 2001). Students working in groups could use the virtual space to coordinate their efforts. In addition, several experts on courthouse design were recruited from around the country to critique students' designs. Results were mixed. Students found that it was easier to coordinate their team in the already substantial in-class time. External critics used the site to comment on designs, but found the lack of dialogue frustrating; they wanted to ask questions of the students. Occasionally, students also were not able to use the critique as they had progressed with their designs before critics could respond. On the positive side, articulating their design on CoWeb caused students to reflect in a different way on their work. Additionally, as critics were able to critique designs, this demonstrated that even a low-bandwidth medium could be useful for discussing designs. For this insight, the project received significant attention from the Architecture community; the American Institute of Architects named CoOL Studio as one of two winning projects nationally in the Education Honors Awards Competition.

Though the CoOL Studio experience was not entirely successful, the Architecture faculty now had a better understanding of the medium and how it might meet their needs. An opportunity to put the medium to use arose quickly. When Georgia Tech switched from quarters to semesters, the introduction classes in several design fields were combined. All

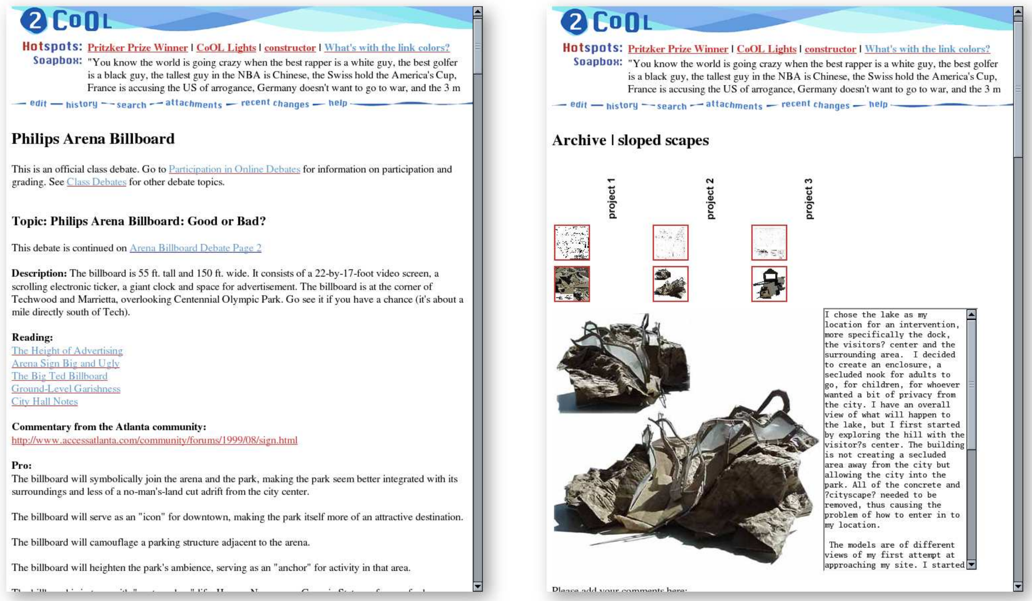

Figure 2: A Discussion and a Gallery in 2CoOL

of a sudden, the studio class approach to the introductory class proved difficult—it just could not scale to serve that many students. Yet, the faculty felt that an active learning style was essential to the education of designers. As such, a new strategy had to be arrived at. CoWeb was to be a fundamental part of the solution.

In fall 1999, 2CoOL came on-line. It served as the on-line space for 171 first-year students for their two-semester introductory class. While the name CoOL stuck, neither the classroom nor web-site was primarily a studio. Because of the class size, the faculty had to revert to a lecture-based class. Yet, they used 2CoOL to add active elements. Because these students would likely share classes again throughout their time at Georgia Tech, it was essential that students formed a community. So, students were given assignments to create personal pages, review movies, and other community building activities. Students also focused on course content. They engaged in large discussions (Craig et al., 2000). The faculty wanted to give students the experience of architecture's collaborative practices, such as design walk-throughs and pin-ups. At the end of the year, they created design galleries (pin-ups) for critics to visit (right part of Figure 2). By the end of the year, 2CoOL contained over 3000 pages.

2CoOL was a success—it served the needs of the learning situation. It also proved to be a challenge, technically. In 1999, CoWeb was completely redesigned to be more flexible and robust. 2CoOL was to be the first real test of the redesign. This class was larger and more intense than any previous CoWeb use. In several ways, the software buckled under the new strain. For the first few months, the server crashed frequently; faculty and students had to learn to accommodate the instability. Many software bugs were found and squashed. The size of the site proved difficult as well; entirely new problems surfaced. For instance, CoWeb allows users to upload files, such as images, to the server. In the original version, all files were added to a central uploads page. This solution was adequate for a small class, but quickly proved problematic in 2CoOL. The page grew so fast as to become immediately unwieldy. To address this, upload directories were added.[2] In addition, the architects wanted new features to serve their needs. Many of these were implemented. For instance, changing the link color based on the age of the page was tried; this innovation proved awkward and was abandoned in later versions. Other changes, such as combining editing with locking control, became standard. Driven by the real needs of 2CoOL, the CoWeb software evolved. It became more powerful, flexible, and robust. In particular, we were able to identify different roles (authors, site designers, purpose agents, etc.) in the process and how those roles could better be supported by the software (Guzdial et al., 2000).

CoOL Studio and 2CoOL were test beds to apply CSCL to architecture education. Faculty worked closely with the developers to apply CoWeb to meet their needs, refining both the use and the medium. Since then, our collaboration with Architecture has transitioned from being a research project to becoming standard practice. The CoOL CoWebs have become a foundation of the introductory class. Since 2CoOL, five more CoWebs with roughly 10,000 pages of content have been created. Currently, the College of Architecture is looking to set up a new server to serve their needs for the foreseeable future.

In English Composition: Learning and Cost

As the software and our understanding of the medium matured, we were ready to more closely analyze CoWeb, its benefits and costs. We were interested in quantifying benefit and cost in a more traditional classroom. For this, we worked with the School of Literature, Communication, and Culture to introduce CoWeb into English composition classes. In comparison to architecture, collaboration was not as core to the field; however, faculty valued collaboration and had previously used other collaborative technologies, such as Web Crossing for chat and discussion boards. Unlike our efforts in architecture, we were not seeking to mature the medium. Instead, we wanted to get a better understanding of the mature medium in a new situation. Could it be used successfully? What were the benefits? What were the costs? To make the application more realistic, we, the developers, took a hands off approach; we did little to train or guide the faculty in use of CoWeb.

This section describes two studies. The first focuses on learning benefit. Two classes are compared: one used CoWeb and the other used a newsgroup-style discussion board. Our results, both qualitative and quantitative, suggest that CoWeb had a significant learning benefit, increasing student performance. Satisfied that there was a clear benefit, we turned our attention to cost in the second study. We discuss both fixed and marginal cost. In particular, we closely examine teacher time as the most significant marginal cost. Our results suggest that CoWeb was a useful medium for engaging students in collaborative learning without incurring disproportionate costs.

[2] While upload directories addressed the immediate problem, they were a mediocre solution. Directories proved unintuitive and awkward to use. Ultimately, page specific upload areas were implemented and directories were deprecated. This story illustrates the iterative development so characteristic of CoWeb's evolution. Through use, a specific problem arose. A solution was developed and applied. Through use, the solution was tested and refined.

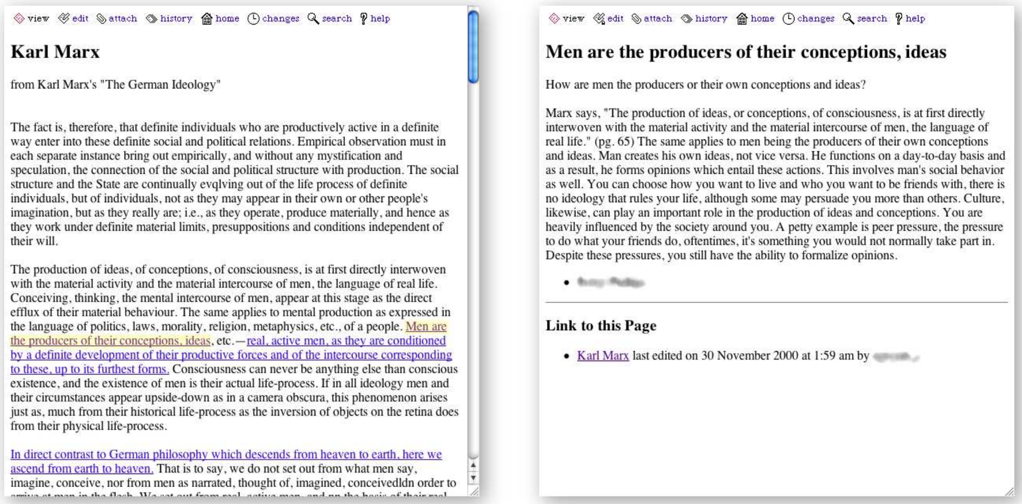

Figure 3: A Close Reading based on the Writings of Karl Marx

Learning

In the first study, we compare two sections of an introductory English composition class, taught by the same instructor.[3] One section (24 students) used CoWeb and the other section (25 students) did not. Both classes completed the same assignments—essays and close readings. The CoWeb section used CoWeb for these assignments. The comparison section completed the essays off-line and the close readings in a threaded-discussion environment. As each section did the same activities, student cost (effort) should be identical. To confirm this, we paid several students in both sections to track their time spent on the class; no notable differences between the groups were observed.

This was the first time this teacher had ever used CoWeb for her class. Even so, she conceived the on-line assignments for the class; she was the first to use close readings in CoWeb. In a close reading, an original source text is annotated with students' comments. In CoWeb, students edit the source text, surround the phrase they wanted to annotate with *s, and save the page. Then, they create a new page to add their annotations to that phrase (Figure 3). In the comparison class, the discussion board was used for close readings: The source text was the original posting and students replied with their annotations. In both sections, students completed close readings based on literature and on student-generated chat sessions in Web Crossing. The instructor imported the chat session into the CoWeb and the newsgroup respectively for the students to annotate.

The results were positive. Through surveys, we found that the CoWeb section had significantly better attitudes toward collaboration than the comparison section (Table 1). In addition, the CoWeb section received higher grades (grade breakdown: 7 As, 10 Bs, 3 Cs, others F or W) than the comparison section (grade breakdown: 19 Bs, 3 Ds, others F or W), which indicates better performance and suggests better learning. In particular, the instructor

[3] The CoWeb section was chosen at random and students did not know a priori which section would use CoWeb, so selection bias was minimized.

Statement	*CoWeb*	*Comparison*	*Difference*
I would rather work independently on assignments than in groups or teams.	2.17	3.19	1.01
I feel working with others on assignments is more helpful than working alone.	4.00	3.25	0.75
When working on team projects, I feel motivated by my sense of responsibility to the group.	4.22	3.31	0.91
I like doing teamwork.	4.11	3.25	0.86
I found it useful to relate my work to that of others.	4.44	3.50	0.94

Table 1: Attitudes toward Collaboration, average scores on a Likert scale where 1 is strongly disagree and 5 is strongly agree. $p < 0.05$ on a two-tailed t-test for all of these statements

noted that the CoWeb section showed more variance; she could assess which students were doing quality work, allowing her to assign a handful of As. In comparison, the performance of the comparison class was uniformly undistinguished.

We recognize that grades are not a precise measure of performance, and they are too large-grained to inform us about where any learning benefit may have come from. As such, twelve students were selected randomly from each section and their work rated by various criteria (Table 2). Five assignments were rated: two close reading assignments based on student-generated chat sessions (rated for the first 6 criteria, which we refer to as chat close readings), two close reading assignments based on literature (rated for the first 10 criteria, referred to as literature close readings), and one formal essay (rated for all 15 criteria). To keep individual bias to a minimum, two raters (one the course instructor, the other a colleague in the same department) rated each assignment on a scale of one to four (four being highest performance). No statistically significant differences were found in their ratings, and all criteria had better than 70% of the ratings identical. In each rating category, the CoWeb section outperformed the comparison section (in most, by a large statistically significant amount).

On average, the students in the CoWeb section did significantly better on writing essays than the comparison section, particularly on issues of vocabulary and essay organization. Several categories show near 1.00 differences in performance; on a scale of one to four, one point of difference indicates a large difference in performance. For instance, on critical vocabulary application, the CoWeb section average is between 2 (chosen when "the student deploys these terms where appropriate in his/her writing, but most are misused") and 3 ("the student deploys most of these terms where appropriate in his/her writing, but occasionally misuses them"), while the comparison section average is between 1 ("the student never successfully deploys these terms where appropriate in his/her writing") and 2.

So overall, we conclude that CoWeb was effective for learning in this study. The performance of the students in the CoWeb section was significantly better by many key subject criteria over the comparison section. At the same time, attitudes towards collaborative

Category	*CoWeb*	*Comparison*	*Difference*
Engagement with Class Material	2.52	1.88	0.64
Foundation for Research	2.49	1.68	0.82
Reflective / Recursive Writing Practices: Authorial Voice	2.30	1.58	0.73
Reflective / Recursive Writing Practices: Reflection and Exploration	2.24	1.49	0.75
Critical Vocabulary: Understanding	2.30	1.54	0.76
Critical Vocabulary: Application	2.28	1.33	0.95
Formation of Critical Questions: Engagement with Topic	2.39	1.94	0.44
Formation of Critical Questions: Quality of Questions / Arguments	2.24	2.21	0.03*
Critical / Close Reading Skills: Analysis	2.29	1.97	0.32*
Critical / Close Reading Skills: Identification of Issues	2.36	2.06	0.31*
Research Skills: Locating Information	3.04	2.54	0.50
Research Skills: Using Information	2.75	2.00	0.75
Identification of Critical Sources	2.75	2.08	0.67
Engagement and Integration of Research Sources	2.71	1.75	0.96
Effective Use of Formal Essay Writing Conventions for Argumentation	2.79	2.21	0.58

Table 2: Writing Performance, average scores on a Likert scale where 1 is lowest performance and 4 is highest performance. $p < 0.05$ on a two-tailed t-test for all except *

learning improved. We speculate that these two factors are not independent; instead, as the use of collaborative learning proves beneficial, more learning will happen, which in turn improves the attitude towards collaboration.

Cost

Now that we have demonstrated a learning benefit, it becomes important to consider cost. What are the costs? Are the benefits worth the cost? Cost breaks down into two categories: fixed and marginal cost (Levin and McEwan, 2001). Fixed costs are infrastructure costs of setting up a CoWeb server. Marginal costs are the costs of maintaining that server and operating classes with CoWeb.

Due to modern computing, fixed infrastructure costs are low. A server was bought for this study, but that server can support at least a dozen classes concurrently over several years. CoWeb is a cross-platform and lightweight server application that can be run on virtually any hardware (in some cases, old 486's), so even an older server can easily support many classes. Student access to Internet-enabled computers is essential for CoWeb use; there was no need to incur any additional infrastructure costs since the campus was already wired for Internet connectivity. Nor is use of that infrastructure markedly increased, considering that students would need similar amounts of time for other applications for the same class (i.e., the discussion board or word processing). At other locations where the infrastructure is not in place, that cost may be prohibitive; however, this infrastructure is becoming very common. The CoWeb software is open-source freeware;[4] thus, there are no software costs.

Administration costs too are negligible. Besides the tracking software (specifically used for gathering study data) and a couple of software upgrades (the CoWeb software was still being evolved), an English professor (not a computer specialist) was able to administer the server without assistance. Across the semester, the total amount of administration time was less than an hour. By far, the dominant cost factor in CoWeb use is instructor time. The instructor for the two sections, using self reporting, averaged about 2.5 hours per week devoted to CoWeb usage; this is quite reasonable as it is about the same amount of time as an office hours session. However, this does not give us a clear idea of how she spent that time or how student usage relates to instructor involvement.

In the term following our learning study, we set up the server to log usage time. We did this for two instructors, teaching the same class (the second English composition course). The first (Instructor 1) was the instructor for the classes in the learning study, and here taught the follow-up course (Class 1: 24 students, with 1 withdrawing). The second (Instructor 2) was the second rater for the performance assessment. This was the first time this instructor used CoWeb, using one CoWeb for three sections of the same class (Class 2: 64 students, with 5 withdrawing). As she was getting used to CoWeb, Instructor 2 still relied on another web environment for the class; in contrast, all on-line activities for Instructor 1 were done with CoWeb.[5] The instructors did different activities with their class and have different styles of using the technology, so this data is a good cross-section of instructional uses. Table 3 summarizes instructor and student time on CoWeb.

[4] It can be downloaded from *http://minnow.cc.gatech.edu/swiki*.

[5] In the future, Instructor 2 plans to only use CoWeb.

	Class 1	*Class 2*
Average Student Time	17.95 hours	8.13 hours
Total Student Time	412.84 hours	484.82 hours
Total Instructor Time	41.30 hours	57.35 hours
Total Student Time / Instructor Time	10.00	8.45

Table 3: Instructor and Student Time using CoWeb

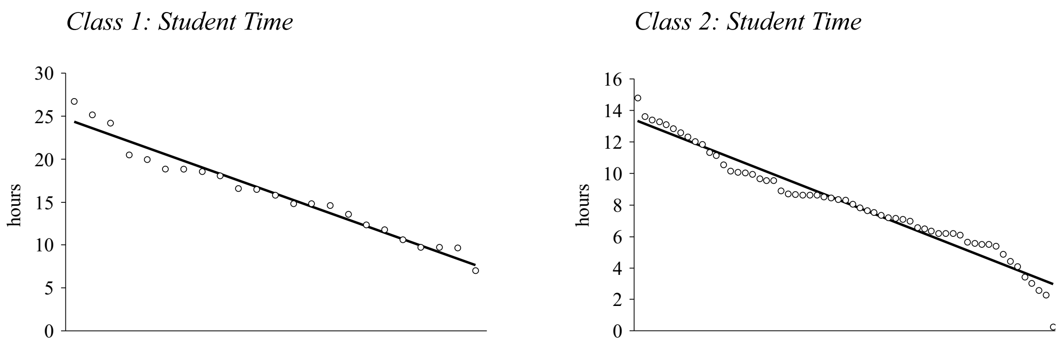

Figure 4: Distributions of Students' CoWeb Usage, from most use to least use

What is most notable is that in both cases the ratio of total time spent by students to total time spent by the instructor is similar (10.00 and 8.45). One way to measure the cost effectiveness of an educational activity is to contrast the ratio of student to instructor time (Levin and McEwan, 2001). By this criterion, lecture is cost effective. For each hour of instructor time, there are *n* hours of total student time (24.00 and 21.33[6] respectively in our case) spent engaged in the learning activity. This number estimate is a bit high, considering it does not include preparation time for the instructor or absenteeism for the students. While lecture scores high marks on efficiency, it loses in learning effectiveness, as student involvement tends to be passive (particularly for large classes where cost efficiency would be high). In contrast, one-on-one tutoring, as may occur during office hours, can be quite active and engaging. Unfortunately, one-on-one tutoring is not economically feasible, with a ratio of 1.00 hour of instructor time to student time. The CoWeb ratios (around 9) on the other hand seem a reasonable compromise of the cost effectiveness of lower instructor time with the learning effectiveness of more active learning (as students construct artifacts).

Unlike lectures that have a high attendance level, time-spent using an educational technology can be highly varied. One scenario could have an exponential drop-off, with only a few students using the technology often. This pattern is often seen in learning activities where participation is voluntary (Bruckman et al., 2002). While the technology might have marked effects on these few students large enough to affect the class average, it probably would not be considered a healthy situation in most schools. What we want to see is that the technology is reaching most if not all students.

To look at the distribution of usage across students, Figure 4 plots student time on

[6] 64 students / 3 sections = 21.33 student class hours per instructor hour

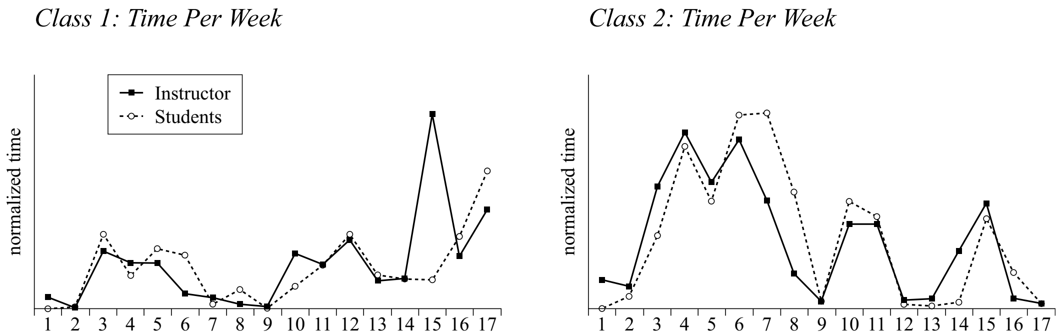

Figure 5: Distribution of Time per Week (note: week 9 is spring break)

CoWeb from most usage to least usage. The vertical axis is the number of hours spent in CoWeb, and the horizontal axis represents different students, ordered in terms of the amount of time they spent in CoWeb. What it shows is that while usage varies quite widely, it does so in a near linear way (for both classes). Also, in both cases, there seems to be a grouping around the class average with only a few doing significantly less or more. This grouping can be seen clearly in the right graph where there is a dip below the best-fit line to the left of the center and a dip above the best-fit line to the right of the center. For an activity, like homework, a roughly linear distribution with a few doing significantly more or less than the average seems acceptable.

Are some activities more cost effective than others (i.e., requiring less instructor time for equal student effort)? If so, efficiency could then be improved by focusing on certain activities and dropping less efficient activities. To test this hypothesis, we recorded student and instructor time on CoWeb over the term; Figure 5 displays the results. The horizontal axis represents week intervals over the course of the term; the vertical axis represents time spent in the CoWeb during that interval. After looking at the data, interviews with the instructors were conducted to find out what activities occurred and how their time was spent.

A couple of conclusions can be drawn from this data. First, almost all of the time, the instructor put in some of the effort before the students; this can be seen particularly well for Instructor 2, where instructor time seems almost shifted a week off the student time. So, a significant proportion of instructor time is spent on setting up the space; this observation was confirmed by both instructors during the interviews. Second, instructor time is closely linked to student time for each assignment. The only exception is week 15 for Instructor 1, where she spent just over 10 hours on CoWeb; this time was mainly spent on grading. Instructor 2 did grading throughout the term. As such, there is no assignment for either instructor that is far more or less efficient. One way to explain this is that the amount of time that instructors and students spend on an assignment is closely related to the point value of the assignment; so, the original hypothesis about more efficient assignments is flawed.

Instructor 2 mainly used CoWeb for one large assignment worth 35 percent of their grade (weeks 2-12). Students worked in small groups (2-3 members) to investigate a decade from 1800-1912. Each group posted a time-line with a minimum of ten significant science or technological innovations or discoveries identified in that decade; each member of the group researched one of these events in depth and wrote a five page paper on it. The purpose of this project was to provide a database of information about science and technology in the 19th century that students could use as background for their final project to create a web-site to understand a 20th century phenomenon in terms of its origins or background in the 19th century. As such, CoWeb served as a research space where students could benefit from the work of their classmates. Although students had to link their final project to the class CoWeb for other students to see, the final projects were required to be traditional web-sites and could not be built in CoWeb. However, the instructor encouraged students to use CoWeb as a way to collaborate on their final project. Most of the use in weeks 13 through 16 is attributable to that voluntary collaboration.

Instructor 1 used CoWeb throughout the term for multiple smaller assignments. Students were required to complete three chat-based and one literature-based close reading assignments. Also, students posted summaries and discussion about the class reading. Instructor 1 also used the space as a way to distribute class readings and communicate deadlines and activities to the students. The largest chunk of student use came during weeks 15 through 17, when they worked on a final project. Like Class 2, the final project for Class 1 was for groups to build a web-site.

Unlike Instructor 2, Instructor 1 allowed students to do their web project entirely in CoWeb; four out of six groups decided to complete their projects entirely in CoWeb. So, students found interaction on CoWeb useful enough to use it instead of traditional website tools, such as Microsoft FrontPage. As students tend to choose the most effective ways to accomplish their goals, this is further evidence of CoWeb's cost effectiveness (this time for students). Furthermore, Instructor 1 commented that the quality of the final projects was higher than previous classes as CoWeb-using students concentrated more on content than on looks. Although the instructor has always stressed content over looks, students creating web-sites tended to spend much of their time on looks. Since most web-site creation tools allow you to mess around easily with looks, it is only natural that students would find this aspect interesting. In contrast, it is almost painful to mess around with looks on CoWeb. Instead of being a detriment in this case, it was an advantage for learning. If CoWeb usage were not seen as cost effective by the students, they would not have used it for their final projects, and the final assignment would not have been as effective for learning. So, it is important that instructor and students see a classroom technology as cost effective. In addition to CoWeb being a useful environment for the final projects, Instructor 1 observed a significant cumulative effect—the CoWeb class was already used to concentrating on content.

For Instructor 1, all class activities, besides office hours and lecture, including grading, were conducted on CoWeb. Considering that lecture time was about 50 hours, roughly 40 hours spent on the class outside of lecture during a semester is quite efficient. The 41 hours observed through system logs also matches closely to Instructor 1's self reported time of 2.5 average hours per week spent on CoWeb for the previous term, where the learning effectiveness was closely examined.

While CoWeb's interface is easy to learn and we (the developers) have produced several guides on how to use it in the classroom, we expect a certain significant cost to be incurred from using a new technology for the first time. As Instructor 1 already used CoWeb before and had taught this course before, her level of efficiency (10.00 total-student-time-to-instructor-time ratio) may have stabilized. In contrast, this was the first time Instructor 2 used CoWeb. As such, her total-student-time-to-instructor-time ratio would be expected to rise (slightly) over time, as she becomes more comfortable with the environment. Also, instructor involvement is highly dependent on teaching style. Instructor 1 views her CoWeb interaction as setting up the space for the students to work and then letting them loose. In contrast, Instructor 2's style is one of tighter control of what occurs in the space; she is actively involved in the running of the activities and likes participating along with the students. This difference in styles might cause Instructor 2's efficiency to be somewhat below Instructor 1's. Even with different styles and uses, CoWeb usage remains cost effective for both instructors.

Reflections

Overall, these results are extremely positive. In the learning study, we demonstrated that one teacher could adopt CoWeb to fit her class. With little guidance, she was able to adapt her curriculum to the new medium. Notably, the medium was particularly suited for her needs—student performance improved measurably. After that initial semester, the teacher grew comfortable enough with the technology to adapt all out-of-class activities to CoWeb. When costs were tracked, it was found that both fixed costs and marginal costs were reasonable for the benefit. Furthermore, we demonstrated that another instructor was able to adopt the technology within a semester and that her time commitment was equally reasonable.

Our work in English composition strengthened our belief in the learning potential of CoWeb. It is likely that others could similarly adopt the technology to benefit their classes. That being said, conditions were good for success. Both instructors are dedicated teachers comfortable with technology and firm believers in collaborative learning. The curriculum too was amiable as the core focus was on writing and the topics addressed lend themselves to open-ended discussion; writing and discussion are core to CoWeb use. It is no wonder that the instructors were able to adopt the medium to serve their curricular needs.

In STEM: Struggling against the Tide

In architecture and English composition, situating CoWeb proved beneficial. Faculty were able to adapt the medium to serve their needs. CoWeb proved to be useful in supporting learning in these situations. Yet, situating CoWeb has not always been successful.

Adoption in Science, Technology, Engineering, and Mathematics (STEM) classes has been overwhelmingly disappointing. Faculty are more reluctant to try CoWeb. Even in the few cases where faculty are enthused and committed, students actively resist: They do not want to collaborate. After continuously struggling against the tide, we gave up on actively situating CoWeb in STEM classes. We shifted our focus to understanding why students did not collaborate. We found three cultural barriers to CoWeb's success. Our evidence leads us to believe that these barriers prohibit the adoption of collaborative media in these fields. The medium cannot be successful while the culture is incompatible.

Anecdotes of Failure

Many senior-level engineering classes require students to use calculus. Yet, Georgia Tech faculty found that most seniors could not properly recall the calculus they had learned previously. The faculty hypothesized that calculus students failed to realize that the mathematics they were learning would be important to their future education. Thus, once their calculus class concluded, they saw no reason to retain those skills. Using CoWeb, a solution was attempted. The Model (short for Mathematical Modeling) CoWeb was designed so that classes in different disciplines (mathematics, chemical engineering, mechanical engineering, etc.) could share the same space. The intention was twofold. First, the engineers would have the resources to remind them of the calculus needed to complete their assignments. Second, the mathematics students could see that the content they were learning would be applicable later.

Results were disappointing. We created a mandatory assignment that required collaboration between two of the classes, a senior-level chemical engineering class and a junior-level mathematics class. The students in chemical engineering created simulations that generated data for the mathematics students to analyze and then provide the results back to the chemical engineers. 40% of the mathematics students accepted a zero on the assignment rather than collaborate with the chemical engineers.

We had a hypothesis that part of the inhibition to participate was a technical one. The contents of these courses involve equations, and equations are difficult to post on the Web. We hypothesized that students avoided using CoWeb because they could not post equations. To test this theory, we designed a browser applet that allowed users to create equations by drag-and-drop. We installed it in the Model CoWeb. Faculty used it and praised it; not a single student in either class even tried it.

The same semester we launched 2CoOL (n=171), we started a CoWeb in a chemical engineering class (n=24). After ten weeks into the semester, the architecture students had generated over 1500 pages, with some discussion pages having over 30 authors (Craig et al., 2000). Not a single chemical-engineering student had made a single posting. In another semester, in a computer science course, only 22 of 340 students ever posted anything.

We have trialed many different CoWeb activities in these classes. Our most successful one was the puzzle activity—the teacher posts a challenging problem and offers extra credit for the solution, for a partial solution, or for leads that results in the solution (Collaborative Software Laboratory, 2000). Approximately 40% of the class voluntarily participated in this activity, which is still a far cry from the 70-100% participation that we see in architecture and English composition.

What's Going on?

These anecdotes paint a stark picture of active resistance to collaboration. These students simply showed no interest in collaborating and, at times, willingly accepted a decrease in their grade rather than collaborate. We do not see that students want to collaborate but are having trouble with the technology or with figuring out how best to collaborate; if that were true, we would expect to see students trying the equation applet and more than 22 of 340 students participating. Rather, we see students actively avoiding collaboration. This is a significant problem, not only because these classes are missing out on the opportunity

Statement	*English*	*Math*	*Chem.E.*
I enjoy using the CoWeb.	3.83	3.48	2.82
I would rather work independently on assignments than in groups or teams.	2.17	2.60	2.41
I feel like working with others on assignments is more helpful than working alone.	4.00	3.64	3.59
I found it useful to relate my work to that of others.	4.44	3.48	3.53

Table 4: Attitudes toward Collaboration in English Composition, Math, and Chemical Engineering, average scores on a Likert scale where 1 is strongly disagree and 5 is strongly agree

Statement	*Chem.E.*	*C.S.*
Posting solutions for comments or questions to the CoWeb is useful.	3.5	3.4
I find the course to take a lot of time outside of class time.	4.2	3.8
I view this field as intensely competitive.	3.9	3.4
I view this class as intensely competitive.	2.4	3.5
Most of the problems in this class have only one correct answer.	3.9	2.3
The CoWeb is primarily an information resource.	3.2	3.1
I print pages from the CoWeb regularly.	2.3	2.2

Table 5: Reasons for Inhibition in Chemical Engineering and Computer Science, average scores on a Likert scale where 1 is strongly disagree and 5 is strongly agree

for better learning, but because the engineering schools accreditation board has mandated collaboration as a critical part of an engineer's education (ABET, 1996).

We have been conducting interviews and questionnaires to try to understand what's going on in these classes. The same term we were studying learning in English composition class, we surveyed the attitudes in a mathematics and a chemical engineering class. Results of an end-of-term survey are summarized in Table 4. We see that the English composition class was more positive about CoWeb and about collaboration in general than the mathematics and chemical engineering classes.

In another study, we used a midterm exam review in a chemical engineering class and in a computer science class; in both classes, there was almost no participation. We used a targeted questionnaire to explore why there was so little participation; the results are summarized in Table 5. In the chemical engineering class (n=24), 90% of the students said that they were aware of the review, and 70% said that they found it useful, but mostly to do on their own. In the computer science class (n=150), 87% of the students said that they were aware of the review, but only 55% found it useful. However, the students generally agree with the statement that "posting solutions for comments or questions to the CoWeb is useful." They just did not do it.

Three Cultural Barriers to Collaboration

Through our research, we discovered three cultural barriers to collaboration in STEM fields: 1) competition and single-answer assignments, 2) the challenge of seeking help, and 3) faculty attitudes and models of collaboration.

Competition and Single-Answer Assignments. STEM students perceived their class or the field as being competitive and demanding (Table 5). Quotes from the targeted questionnaire on why students did not participate in the midterm exam review bolstered this claim. Students responded that they "didn't want to get railed" and that "with the curve, it is better when your peers do badly." "Since it is a curved class, most people don't want others to do well." Even when the instructor made it clear that the class was not curved, students still perceived it to be. "It has to be curved. He can't flunk us all." Instead of collaborating to improve their performance, students counted on a beneficial grade curve.

Students in engineering and mathematics, additionally, tended to see their homework as having only one correct answer (Table 5), even when faculty stressed that this was not the case. Again, it was just the students' perception. If only one correct answer exists (or is perceived to exist) and the class is highly competitive and/or curved, it is rational for students not to collaborate or help others. It is in a student's best interest not to participate. Research on collaborative learning informs us that the perception of single-answer assignments is a hindrance to collaboration. Cohen (1994) in her review of the literature on collaborative learning found that open-ended ill-structured problems tend to encourage productive group learning; if the students perceive that there is only one answer, there is not as much need for the group.

The Challenge of Seeking Help. Educational psychology has documented a paradox in student behavior when choosing to seek help: If a student is confused, he may not want to seek help, perhaps to avoid admitting the confusion. This condition is called *learned helplessness* (Bruer, 1993). Seeking and receiving help does lead to achievement, but students have to seek the help (Webb and Palincsar, 1996). Quotes from the targeted questionnaire support the belief that the students may have felt that they were so confused that they could not ask for help:

"I haven't posted about questions because I am confident that my answers are wrong." A wrong answer, however, could invite help or lead someone else to a correct answer. "I thought I was the only one having problems understanding what was asked in the exam." Given the general confusion in the class, this student was not the only one struggling. "Who am I to post answers?"

Students also felt that if they asked questions, they would be punished in the competitive atmosphere: "What was I suppose to do with it? Those who answered questions were severely criticized by [the teacher]." "The overall environment for [this class] isn't very help-oriented."

Faculty Attitudes and Models of Collaboration. As these quotes demonstrate, some faculty were a detriment to a collaborative atmosphere. We have witnessed faculty, who were persuaded by their students to start a CoWeb, proceed to lock down the vast majority of pages, thereby effectively stifling collaboration. One civil engineering faculty, upon learning about our results, responded, "but, undergraduate students should have only single-answer problems! Design comes much later!" When posed the issue about ill-structured problems better supporting collaboration, he replied that he did not believe that collaboration was important. He was aware of the ABET (1996) mandate, but chose to ignore it. We have had similar responses from other faculty and teaching assistants with whom we have spoken. If undergraduate learning is about learning facts and skills, then where is the role for collaboration?

If faculty are not supportive of collaboration, they may not convey to students what collaboration is about or how or why they should collaborate. Or, even if the faculty are supportive, a traditional lecture-style class may not provide students with the models for what they are supposed to do in a collaborative learning situation. In interviews, engineering and computer science students informed us that they did not collaborate in CoWeb because they simply did not know what to do there. The students had no models for how to collaborate, nor how to learn collaboratively (at least, in these classes).

Offering the Faculty an Opportunity to Change

After these studies, we realized that the best opportunity for change was to directly address the faculty who might be interested in using CoWeb. In Spring 2001, we offered a workshop to Georgia Tech faculty who wanted to use CoWeb. During a two-hour lunchtime session, we led a dozen faculty through using CoWeb for themselves (each had their own station). We had three faculty talk about how they used it. We also offered the faculty support documentation, including a copy of the catalog (Collaborative Software Laboratory, 2000) of the activities that teachers had invented in their own courses. Each of the faculty used CoWeb during the workshop, and all expressed satisfaction (on an exit survey) that it was usable for their courses and by them. At end of Summer 2001, we followed up with each of the faculty and offered them additional support, including offers to create and host CoWebs for them on our own servers.

In November 2001, we followed up with the faculty who took our workshop. Only one faculty member (from Psychology) had started using CoWeb. The rest, including mathematics and engineering faculty, had not adopted it. We surveyed all of the faculty. The common explanation was a lack of time to explore new options in their classes. We used the same survey with a group of faculty actively using CoWeb. We found that all those teachers who were actively using CoWeb were already using some form of collaborative learning in their courses. For the teachers already looking for a mechanism to encourage collaborative learning, CoWeb met a need and was thus cost effective. For the others, the barrier to entry was too high; the perceived benefit was not worth the perceived cost. Given our results, it's hard to disagree with that conclusion.

Disappointed, we have given up actively trying to situate CoWeb in STEM classes; we have switched to a passive strategy of "let them come to us."[7] The cultural barriers are just too high for casual success. In the next section, we detail an exception. We summarize

[7]Since then, a few faculty in STEM fields (in particular, chemical engineering and biomedical engineering) have adopted CoWeb. In those cases, the faculty had a specific need for CoWeb. We have helped them set up CoWebs for their classes.

how one STEM class was designed, albeit at substantial cost, to encourage a collaborative culture and the substantial role that CoWeb played in realizing this design.

In Media Computation: Designing Culture

This section details CoWeb's role in the Introduction to Media Computation (MediaComp) course at Georgia Tech. This introductory computing course is aimed at students in liberal arts, management, and architecture. Unlike our other efforts in STEM, we were not simply situating CoWeb into an already established class. Instead, the course was designed around CoWeb: Our goal was to create a collaborative context in a computer science course. The course was designed from its inception to encourage collaboration, with explicit encouragement to share artifacts created in the class (Guzdial, 2003). Creating this class cost substantial time and effort; however, there was a payoff: The results were much more fruitful than we saw in the other STEM courses.

Framing the Problem

Like in the STEM courses described in the previous section, we are working against significant cultural barriers in the field and in the classroom. The experience of introductory computing courses for most students are solitary affairs (AAUW, 2000). While there are efforts (Nagappan et al., 2003; McDowell et al., 2002) to make introductory computing courses much more collaborative, these are still quite novel and are not frequently practiced.

Over the last five years, the overall number of students studying computer science in the United States has dropped rapidly—a 50% or more decline at many schools (Vegso, 2005), but this is just the steepest part of the decline. Overall, interest among incoming U.S. freshmen in computing has dropped 70% since the late 1980's (Vegso, 2005). The problem of declining interest in computing is particularly severe among women and minorities where there has been a decline in numbers of students and percentage of majors for the last 30 years (AAAS, 2005; Margolis and Fisher, 2002).

Several studies (AAAS, 2005; AAUW, 2000; Margolis and Fisher, 2002; Pfleeger et al., 2001) have explored why there is such a declining interest in computing in the U.S., and they have all come to similar conclusions. Computing courses and computing as a field are perceived as being asocial, lacking in opportunities for creativity, and boring with too much attention on technical details. Most significantly, what appears in the introductory computing courses is seen as irrelevant—students complain that they see no connection between the topics and assignments in the introductory course and the problems of real people (Margolis and Fisher, 2002; AAUW, 2000).

A measurable manifestation of the failure of introductory computing courses to engage students is the withdrawal-failure-or-D (WFD) rate—the percentage of students who withdraw from the course or earn a D or F (not passing for most majors). There are no national or international surveys suggesting what overall WFD rates look like, but estimates suggest 30–50% WFD rates (Roumani, 2002), and empirical studies of WFD rates fall in that range (Nagappan et al., 2003). Non-CS-majors' WFD rates seem to be significantly higher, e.g., 60% at least in some studies (Nagappan et al., 2003).

It was against this backdrop that Georgia Tech began an initiative to improve the success rate in its own introductory computing course. Every student at Georgia Tech is required to take an introductory course in computing. From Fall 2000 through Fall 2002, only one course met that requirement. The average WFD rate during those years was 28.2%. The inverse of that number is the success rate—the percentage of students who complete the course and earn an A, B, or C. Thus, the average success rate was 71.8%. While this was well within the norms of the field, we wanted to improve it.

The decision was made to create contextualized computing courses. Rather than a generalized abstract course on computing, two new additional classes would draw upon the students' majors to find relevant examples and homework assignments. One course focused on engineering uses of computing concepts and programming. The second course, the MediaComp course, was designed for students in liberal arts, management, and architecture.

The MediaComp course introduced the same computing concepts as in a traditional introductory course as defined in recommendations from national professional organizations (ACM/IEEE, 2001), but using media-related examples. To learn about iteration, students would write programs to iterate across pixels to negate or gray-scale an image or increase the volume of a sound. To learn about ranges of indices, students would remove red-eye from an image, or splice a sound. Students would also create and manipulate web pages, explore digital video special effects, and combine these (e.g., writing a program to parse headlines from a news web-site and then creating a ticker-tape movie from these headlines).

The MediaComp course sought to address the needs of the students who traditionally are least likely to succeed in a computing course. Compared to engineering and computing majors, the students were less likely to have the strong background in mathematics and prior computing experience, correlated with success in computing (Wilson and Shrock, 2001). We, the developers, chose to address the concerns identified in the literature as problematic in introductory computing courses. We wanted to create a course that was perceived as being relevant, as offering opportunities for creativity, and (most significantly for this article) as offering a social and supportive community (Forte and Guzdial, 2004; Guzdial and Forte, 2005).

Constructing a Collaborative Context

The choice of media computation as the context for the course explicitly addresses the findings of the studies identifying challenges for computing students. For the professions most likely to be pursued by these students, the computer is less a tool of calculation as it is a tool for communication (Forte and Guzdial, 2004). Our expectation is that professionals who graduate from liberal arts, management, and architecture programs tend to use the computer via applications such as Microsoft Powerpoint, word processors, and image-manipulation programs (e.g., Adobe Photoshop). Even the use of calculation tools among this group (e.g., spreadsheets such as Microsoft Excel) are often used for the creation of charts and presentation graphics. Our expectations were supported by our discussions with faculty in these disciplines as part of our design process (Guzdial and Forte, 2005).

Media computation as a context is therefore relevant for this audience, to address the issue of irrelevancy in introductory computing courses. It is authentic (Shaffer and Resnick,

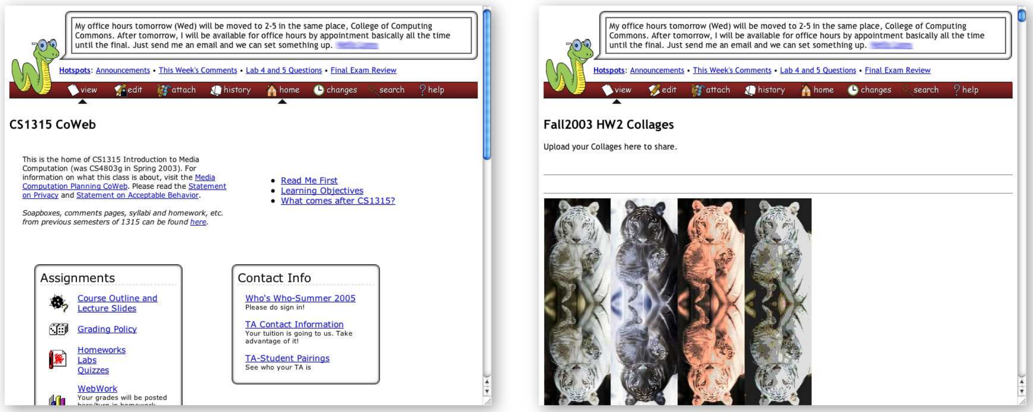

Figure 6: The Home Page and a Gallery on the MediaComp CoWeb

1999) in the sense of being the kind of computation that experts in the students' chosen domain use. In addition, homework assignments in this context could easily offer the opportunity for creative expression.

The use of the CoWeb in the course would allow us to embed the relevant and creative activity in a social setting. We chose to adopt uses of the CoWeb drawn from our prior observations of and experience with different kinds of classes. As in many course CoWebs, we created a "Who's Who" page on which students would create their own home-page within the course and introduce themselves (Guzdial et al., 2001). We created pages for homework questions and answers, and for exam reviews where students could post their answers to sample questions from past examinations.

In particular, we adapted several of the elements that were successful in architecture. A "soapbox" (the snake's speech bubble in Figure 6) enabled students to post messages to appear at the top of every page, creating a course-wide forum for news, ideas, and community building. We adopted design galleries in which students could share their work (right half of Figure 6). If we were successful in creating assignments that encouraged creative expression, on-line galleries could provide the opportunity to share that creative output, and thus provide a significant social setting for the course.

We did not attempt to measure the cost of creating the course, but it was significant. The cost is hard to estimate because it was spread across many people over a long period of time. The idea for the use of the media computation context in introductory computing started among the authors and our collaborators in 2001. A computer music implementation course was offered as a special topics course for upper-level undergraduate students in order to start exploring assignments for an introductory course (Guzdial and Greenlee, 2002). The decision to offer the MediaComp course was made in Spring 2002, with the first offering in Spring 2003.

During the intervening year, the course had to be assembled from scratch, since no similar course had been offered anywhere previously. The challenge of creating the course was not just its definition, but creating enough materials to support a large course. The first offering of the course had 120 students, and the average of the semesters since is around 300 students.

Our design process involved a significant cost. We involved stakeholders in several ways (Guzdial and Forte, 2005): Students were polled from freshmen and sophomore classes across campus, to ask about their interests and goals for a computing class. A faculty advisory board was formed with representatives from the liberal arts, management, architecture, and interested faculty in engineering and sciences. These faculty helped answer questions about their majors and their disciplines, and provided feedback on the evolving course design. A month before the class started, a one day workshop was held where interested faculty could review materials and try some of the assignments that students would experience in the course.

The development process for the course materials involved literally dozens of students. For example, a programming environment had to be constructed for the course. The programming language with which we were teaching, Python,[8] was not commonly used for introductory computing students, but met the criteria of our faculty advisory board and student informants (Guzdial and Forte, 2005). Existing Python programming environments were designed for professional programmers and were inappropriate for novices (Guzdial, 2004b). A team of undergraduate students built the first version of JES (Jython Environment for Students) for their senior design project in Summer 2002, and then other students evolved the environment (e.g., improved error messages, added additional features) over the last three years.

The largest cost in creating the MediaComp course was creating the course content: student activities, lectures, and examples. Undergraduate students who succeeded in graphics and digital video special effects courses were recruited to help create simple versions of their programs in Python to be used as examples. A textbook (Guzdial, 2004a) was created to support the course. Finally, student homework assignments had to be defined. The assignments were created to offer open-ended opportunities for creativity. For example, an early assignment (at week four or six in the fifteen week semester) asks students to create a collage where some small picture is composed into a collage canvas at least four times, where one time is just a copy, but the other three involve some kind of image manipulation: color modification, cropping, or scaling (right half of Figure 6). Students are then invited to share their collages in the gallery page, to share with their peer students.

Results

In the initial offering of the course, we made a significant effort to measure the success of the design decisions that we had made (Rich et al., 2004; Forte and Guzdial, 2004). We conducted surveys at the start, midterm, and completion of the course. We also conducted interviews with the women in the course at two points during the semester.

From the start, we achieved the goal of improved success in the course. Table 6 presents the success rates of the first year of the course. Compared to the 71.2% of the original course and considering that the students were neither engineering, nor computing majors, the success rates have been quite high.

As in 2CoOL, the use of the CoWeb was integral from the very beginning. Literally hundreds of pages were created each semester, with many artifacts created from homework

[8]Specifically, the Jython dialect: *http://www.jython.org*.

Term	*Enrollment*	*Success Rate*
Spring 2003	120	90.0%
Fall 2003	303	86.5%
Spring 2004	395	86.9%

Table 6: Enrollment and Success Rates in MediaComp

assignments (e.g., visual collages, audio collages, digitally created animations) being shared on gallery pages. Students told us of the value of the galleries in interviews:

> It's nice to see other people, like what they did with it.
>
> I don't ever look at it [the homework gallery] until after I'm done. I have a thing about not wanting to copy someone else's ideas. I just wish I had more time to play around with that and make neat effects. But JES will be on my computer forever, so... the nice thing about this class is that you could go as deep into the homework as you wanted. So, I'd turn it [the homework assignment] in, and then me and my roommate would do more after to see what we could do with it.

Female students whom we targeted in interviews and surveys did tell us that they found the course to be relevant, creative, and social:

> [On relevance] I dreaded computer science, but *all* of the topics thus far have been applicable to my future career (and personal) plans—there isn't anything I don't like about this class.
>
> [On opportunities for creativity and social interaction] Actually, I think [collaboration] is one of the best things about this class. My roommate and I abided by all the rules,[9] but we took full advantage of the collaboration. It was more just the ideas bouncing off each other. I don't think this class would have been as much fun if I wasn't able to collaborate.
>
> [On CoWeb use] It's not just about the class... People talk about anything, it's a little bit more friendly than just "here's your assignment."

The challenge of the course design was to create a classroom context for collaboration in a traditionally problematic domain, computer science. Perhaps a sign of the success of the course design was that survey results suggested that students thought of MediaComp as something different than computer science. In the Fall 2003 final course survey, we asked students about their interest in taking additional computer science and additional media computation classes; their Likert scale responses are listed in Table 7. Considering that

[9] The student is indicating that her collaboration was within the rules, avoiding cheating such as simply copying the roommate's code. An impetus for creating contextualized classes was that there was a major cheating scandal in the old class. As such, the difference between collaboration and cheating was still on students' mind.

I'm interested in taking further classes in...	*Strongly Disagree*	*Disagree*	*Neutral*	*Agree*	*Strongly Agree*
...Computer Science	15.2%	37.7%	23.8%	19.3%	4.1%
...Media Computation	11.2%	22.9%	23.3%	26.5%	16.1%

Table 7: Interest in further Computer Science or Media Computation Courses

MediaComp is computer science, one would expect responses to computer science to be at least as high. Students were more excited about media computation.

Other schools are now adopting our course. We have asked the same questions about additional courses at other institutions. In general, the results are the same—that a media computation approach successfully creates a new and more welcoming context, but that context is not necessarily seen as computer science. The process we describe here was successful in creating an entirely new context in which collaboration is successfully engendered.

Reflections

It has been eight years since we first conceived of applying WikiWikiWeb to learning situations. Since then, we have designed our own Wiki implementation, CoWeb, to serve learning purposes and supported its adoption in a wide variety of authentic contexts. We have seen CoWeb support collaborative learning in a variety of disciplines. Wiki has proved to be a useful medium: Teacher and students were able to go beyond adopting the medium to inventing new uses to serve their needs (Guzdial et al., 2001).

Results

We have seen success, particularly in fields and classes that value open-ended discussion and reflection through writing. In architecture, the medium enables a large class to still be grounded in active learning and to maintain a community feel. In English composition, teachers were able to adopt the technology to benefit learning at low cost. Students were able to concentrate on content (rather than looks) to achieve better results.

Yet, we have also seen failure. In STEM courses, significant barriers prevented the medium from being successfully adopted. These barriers are not a product of the medium, but a product of the cultural context. Thus, no matter how much we improved CoWeb, the barriers still existed. Many students never even tried CoWeb. To summarize these findings, *culture trumps medium*. If a culture is incompatible with a medium, the medium will fail. Yet, we have also seen that a medium can be a powerful tool for changing culture. In MediaComp, CoWeb served as an essential part of a design to change the culture of an introductory computer science class.

Approach

In those eight years, we have also realized that our research approach differs from previously articulated, established approaches. While our work is design-based research, we do not tightly control the context of use, as in a design experiment. Instead, we try to support others in situating a new medium to serve their needs. This is a messy, situated enterprise. Yet, it is necessary to properly understand the potential of the Wiki. In addition to learning benefit, costs are important. If a medium is to be used, that use must be practical as well. Thus, we situate our approach as a scholarship of application. There is a need for this kind of research. By detailing our work in this article, we aim to demonstrate the importance of this mode of scholarship to the learning sciences.

Acknowledgments

Funding for this project was provided by the Mellon Foundation and three grants from the National Science Foundation (REC-9814770, DUE, and CCLI Educational Innovations). Thanks to our faculty collaborators: Jim Greenlee, Lissa Holloway-Attaway, Sabir Khan, Pete Ludovice, Tom Morley, Matthew Realff, Brandy Walker, and Craig Zimring. Thanks to the students that contributed to CoWeb development and research: Karen Carroll, David Craig, Andrea Forte, Joshua Gargus, Colleen Kehoe, Bolot Kerimbaev, Akbar Ladak, Kayt Sukel, and Allison Elliott Tew.

References

AAAS (2005). *Preparing Women and Minorities for the IT Workforce*. American Association for the Advancement of Science, Washington DC.

AAUW (2000). *Tech-Savvy: Educating Girls in the New Computer Age*. American Association of University Women Education Foundation, New York.

ABET (1996). Engineering criteria 2000: Criteria for accrediting programs in engineering in the United States. *ASEE Prism*, March:41–42.

ACM/IEEE (2001). Computing curriculum 2001. *http://www.acm.org/sigcse/cc2001*.

Barab, S. and Squire, K. (2004). Design-based research: Putting a stake in the ground. *The Journal of the Learning Sciences*, 13(1):1–14.

Berners-Lee, T. (1999). *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by its Inventor*. HarperSanFranciso.

Bolter, J. D. and Grusin, R. (1999). *Remediation: Understanding New Media*. MIT Press, Cambridge, MA.

Boyer, E. L. (1990). *Scholarship Reconsidered: Priorities of the Professoriate*. Jossey-Bass Books.

Brown, A. L. (1992). Design experiments: Theoretical and methodological challenges in creating complex interventions in classroom settings. *The Journal of the Learning Sciences*, 2(2):141–178.

Bruckman, A., Jensen, C., and DeBonte, A. (2002). Gender and programming achievement in a cscl environment. In Stahl, G., editor, *Proceedings of CSCL 2002*, pages 119–127.

Bruer, J. (1993). *Schools for Thought: A Science of Learning in the Classroom*. MIT Press, Cambridge, MA.

Cohen, E. G. (1994). Restructuring the classroom: Conditions for productive small groups. *Review of Educational Research*, 64(1):1–35.

Collaborative Software Laboratory (2000). A catalog of coweb uses. Technical Report GIT-GVU-00-19, Georgia Tech GVU Center.

Collins, A. (1992). Towards a design science of education. In Scanlon, E. and O'Shea, T., editors, *New Directions in Educational Technology*, pages 15–22. Springer-Verlag, New York.

Craig, D., ul Haq, S., Khan, S., Zimring, C., Kehoe, C., Rick, J., and Guzdial, M. (2000). Using an unstructured collaboration tool to support peer interaction in large college classes. In Fishman, B. J. and O'Conner-Divelbiss, S. F., editors, *Proceedings of ICLS 2000*, pages 178–184. Lawrence Erlbaum Associates.

Craig, D. L. and Zimring, C. (2000). Supporting collaborative design groups as design communities. *Design Studies*, 20(2):187–204.

diSessa, A. A. (2000). *Changing Minds: Computers, Learning, and Literacy*. The MIT Press, Cambridge, MA.

Fishman, B., Marx, R. W., Blumenfeld, P., Krajcik, J., and Soloway, E. (2004). Creating a framework for research on systemic technology innovations. *The Journal of the Learning Sciences*, 13(1):43–76.

Forte, A. and Guzdial, M. (2004). Computers for communication, not calculation: Media as a motivation and context for learning. In *Hawaii International Conference on System Sciences*.

Guzdial, M. (2003). A media computation course for non-majors. In *Proceedings of the Innovation and Technology in Computer Science Education (ITiCSE) 2003 Conference*, pages In–Press, New York. ACM, ACM.

Guzdial, M. (2004a). *Introduction to Media Computation: A Multimedia Cookbook in Python*. Prentice-Hall.

Guzdial, M. (2004b). Programming environments for novices. In Fincher, S. A. and Petre, M., editors, *Computer Science Education Research*, pages 127–154. Taylor & Francis Group, PLC, London.

Guzdial, M. and Forte, A. (2005). Motivation and non-majors in computer science: Identifying discrete audiences for introductory courses. *IEEE Transactions on Education*, 48(2):248–253.

Guzdial, M. and Greenlee, J. (2002). A computer music implementation course using active essays. Gvu technical report tr #git-gvu-02-08, GVU Center, College of Computing, Georgia Institute of Technology.

Guzdial, M., Rick, J., and Kehoe, C. (2001). Beyond adoption to invention: Teacher-created collaborative activities in higher education. *The Journal of the Learning Sciences*, 10(3):265–279.

Guzdial, M., Rick, J., and Kerimbaev, B. (2000). Recognizing and supporting roles in cscw. In *Proceedings of CSCW 2000*, pages 261–268. ACM Press.

Hoadley, C. P. (2002). Creating context: Design-based rearch in creating and understanding cscl. In Stahl, G., editor, *Proceedings of CSCL 2002*, pages 453–462.

Kay, A. and Goldberg, A. (1977). Personal dynamic media. *IEEE Computer*, March 1977:31–41.

Kehoe, C. (2001). *Supporting Critical Design Dialog*. PhD thesis, Georgia Institute of Technology, College of Computing, Atlanta, GA.

Koschmann, T., editor (1996). *CSCL: Theory and Practice of an Emerging Paradigm*. Lawrence Erlbaum Associates, Mahwah, NJ.

Lave, J. (1997). What's special about experiments as contexts for thinking. In Cole, M., Engeström, Y., and Vasquez, O., editors, *Mind, culture, and activity: seminal papers from the Laboratory of Comparative Human Cognition*, pages 57–69. Cambridge University Press.

Lave, J. and Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, Cambridge, UK.

Leuf, B. and Cunningham, W. (2001). *The Wiki Way: Quick Collaboration on the Web*. Addison-Wesley.

Levin, H. M. and McEwan, P. J. (2001). *Cost-Effectiveness Analysis*. Sage Publications, Inc., second edition.

Margolis, J. and Fisher, A. (2002). *Unlocking the Clubhouse: Women in Computing*. MIT Press, Cambridge, MA.

McCandliss, B. D., Kalchman, M., and Bryant, P. (2003). Design experiments and laboratory approaches to learning: Steps toward collaborative exchange. *Educational Researcher*, 32(1):14–16.

McDowell, C., Bullock, H., Fernald, J., and Werner, L. (2002). The effects of pair-programming on performance in an introductory programming course. In Knox, D., editor, *The Proceedings of the Thirty-third SIGCSE Technical Symposium on Computer Science Education, 2002*, pages 38–42. ACM, New York. Pair-programming (social) reduces attrition rates.

McLuhan, H. M. (1962). *The Gutenberg Galaxy: The Making of Typographic Man*. University of Toronto Press.

McLuhan, H. M. (1964). *Understanding Media: The Extensions of Man*. The MIT Press, Cambridge, MA.

Meyrowitz, J. (1985). *No Sense of Place: The Impact of Electronic Media on Social Behavior*. Oxford University Press, New York, NY.

Mortensen, T. and Walker, J. (2002). Blogging thoughts: personal publication as an online research tool. In Morrison, A., editor, *Researching ICTs in Context*, pages 249–279. InterMedia Report.

Nagappan, N., Williams, L., Ferzil, M., Wiebe, E., Yang, K., Miller, C., and Balik, S. (2003). Improving the cs1 experience with pair programming. In Joyce, D. and Knox, D., editors, *Twenty-fourth SIGCSE Technical Symposium on Computer Science Education*, pages 359–362, New York, NY. ACM.

Papert, S. (1993). *Mindstorms: Children, Computers and Powerful Ideas*. Basic Books, New York, NY, second edition.

Pfleeger, S. L., Teller, P., Castaneda, S. E., Wilson, M., and Lindley, R. (2001). Increasing the enrollment of women in computer science. In McCauley, R. and Gersting, J., editors, *The Proceedings of the Thirty-second SIGCSE Technical Symposium on Computer Science Education*, pages 386–387. ACM Press, New York.

Radick, G. (2003). R. l. garner and the rise of the edison phonograph in evolutionary philology. In Gitelman, L. and Pingree, G. B., editors, *New Media, 1740-1915*, pages 175–206. The MIT Press, Cambridge, MA.

Resnick, M., Bruckman, A., and Martin, F. (1996). Pianos not stereos: Creating computational construction kits. *Interactions*, 3(6):41–50.

Rich, L., Perry, H., and Guzdial, M. (2004). A cs1 course designed to address interests of women. In *Proceedings of the ACM SIGCSE Conference*, pages 190–194, Norfolk, VA.

Rick, J. and Lamberty, K. K. (2005). Medium-based design: Extending a medium to create and exploratory learning environment. *Interactive Learning Environments*.

Roumani, H. (2002). Design guidelines for the lab component of objects-first cs1. In Knox, D., editor, *The Proceedings of the Thirty-third SIGCSE Technical Symposium on Computer Science Education, 2002*, pages 222–226. ACM, New York.

Scardamalia, M. and Bereiter, C. (1991). Higher levels of agency for children in knowledge building: A challenge for the design of new knowlege media. *The Journal of the Learning Sciences*, 1(1):37–68.

Schön, D. A. (1987). *Educating the Reflective Practitioner*. Jossey-Bass Inc.

Shaffer, D. W. and Resnick, M. (1999). “thick” authenticity: New media and authentic learning. *Journal of Interactive Learning Research*, 10(2):195–215.

Songer, N. B. (1998). Beyond "real school" obstacles: Ten essentials for adoption of technology-rich programs in urban middle schools. In Bruckman, A. S., Guzdial, M., Kolodner, J. L., and Ram, A., editors, *Proceedings of ICLS 1998*, pages 270–276. AACE.

Stone, A. R. (1998). *The War of Desire and Technology at the Close of the Mechanical Age*. The MIT Press, Cambridge, MA.

Turkle, S. (1984). *The Second Self: Computers and the Human Spirit*. Simon and Schuster, New York, New York.

Turkle, S. (1995). *Life on the Screen: Identity in the Age of the Internet*. Simon and Schuster, New York, New York.

Vegso, J. (2005). Interest in cs as a major drops among incoming freshmen. *Computing Research News*, 17(3).

Webb, N. M. and Palincsar, A. S. (1996). Group processes in the classroom. In Berliner, D. C. and Calfee, R. C., editors, *Handbook of Educational Psychology*, pages 841–873. Simon & Schuster Macmillan.

Wilson, B. C. and Shrock, S. (2001). Contributing to success in an introductory computer science course: A study of twelve factors. In McCauley, R. and Gersting, J., editors, *The Proceedings of the Thirty-second SIGCSE Technical Symposium on Computer Science Education*, pages 184–188. ACM, New York.

Yin, R. K. (2003). *Case Study Research: Design and Methods*. SAGE Publications, third edition.

Zimring, C., Khan, S., Craig, D., ul Haq, S., and Guzdial, M. (2001). Cool studio: Using simple tools to expand the discursive space of the design studio. *Automation in Construction*, 10(6):675–685.