Educating Novice Developers of Notification Systems: Targeting User-Goals with a Conceptual Framework

C. M. Chewar and D. Scott McCrickard
Department of Computer Science
Virginia Polytechnic Institute and State University
United States
{cchewar, mccricks}@cs.vt.edu

Abstract: We approach the challenge of educating undergraduate and graduate students in human-computer interaction (HCI) of notification systems. Notification systems are software and physical interfaces that deliver information of interest to users engaged in other activities. Often, such systems support educational activities and facilitate groupwork, but their range of use extends far beyond these domains and design challenges hinge on effective multitasking and attention management approaches. We have created a framework for developing notification systems that helps HCI researchers and practitioners appreciate these challenges, articulate their system goals in a common representation, and leverage development resources. Through this framework, resources that include information and interaction design guidelines, existing system examples, specific prototyping techniques, and evaluation approaches are accessible to a system developer. We present the outcomes of an instructional approach founded on various educational theories that integrates the framework to present key phases of system design and evaluation.

Designing Notification Systems: An HCI Challenge

We are engaged in an on-going project to improve undergraduate and graduate education of human-computer interaction (HCI) for developing notification systems. Notification systems are software and physical interfaces that deliver information of interest to users that are engaged in other activities. Often, such systems support educational activities and facilitate groupwork, but their range of use extends far beyond these domains. For example, the Irwin notification system presents a compact, yet detailed view of activities within information resources, such as Usenet newsgroups, listservs, and discussion forums (McCrickard 1999)ó certainly providing a valuable mechanism for monitoring of distance education group interaction. While Irwin and many other notification systems reside in a small corner of a desktop screen, other notification systems have been developed on large screen display or wall mounted platforms and as physical objects endowed with dynamic, computationally modifiable features. The notification system for the Virtual School project is an example of a large screen display that supports classroom educationó using a timeline that reflects current progress toward group activities, the display enables students and teachers to gain awareness about each otherís actions and react appropriately (Carroll et al *in press*). Socially translucent notification systems can also enhance instruction by conveying visualization of group member activity patterns (Erickson et al., 2002).

Effective and elegant design of these systems is a growing concern within the HCI community, largely due to peopleís willingness and desire to monitor multiple sources of information during daily tasks and the advances in data delivery technologies. To produce effective interfaces, developers must clearly understand the user priorities related to willingness to be interrupted from primary tasks (referred to henceforth as *interruption*), desire to react to notification content (*reaction*), and achievement goals for long-term *comprehension* of information (McCrickard & Chewar, 2003). We have argued that design choices for interface features of notification systems should be guided by the effect on three critical parametersó interruption, reaction, and comprehension. While our research is focused on improving the design-science and usability of these systems, we are also interested in teaching computer science students to appreciate these challenges and inspire them to address such problems in their own research.

However, we have found that addressing these concerns requires mastery of a wealth of human factors and experimental psychology research results, such as those found in (Wickens & Hollands, 2000). Slight differences in information presentation, such as color or size of interface elements and animation strategies, can have significant and unwanted effects on distraction to the primary task or interpretation and processing of the notification content. Unfortunately, most computer science students receive little or no formal instruction on design considerations or in related academic disciplinesó especially those specifically addressing attentional and multitasking concerns. While style guides and design templates help guide novice interface designers through standard interface design and usability engineering challenges, similar tools do not exist for notification systems. To fill this gap, we have developed a conceptual framework for relating combinations of user goals (in terms of the three critical parameters) to resources that include information and interaction design guidelines, existing system examples, specific prototyping techniques, and evaluation tools and approaches. We describe this framework and then explain how we integrated it into three different versions of semester-long instructional programs. Following this is an analysis of the outcomes achieved by using this approach, as well as a discussion of future work.

Conceptualizing User Goals for Notification Systems

To improve our ability to practice ìgood HCIî during the design and evaluation of notification systems, we have developed a conceptual framework that generalizes user goals for the systems and provides association to development resources. While this paper focuses on how we used the framework to support an HCI education process, in addition we hope that this conceptualization can also be valuable to the ED-MEDIA community for designing and evaluating multitasking or notification systems used to support instruction. In this section, we provide an overview of the *IRC framework*, although an extended introduction to this framework may be found in (McCrickard & Chewar, 2003). Our description begins with the basic mechanics of the framework, to include processes of generalizing user goals and visually representing a systemís *design model* (a designerís understanding of user requirements and usage context). Representing combinations of user goals visually helps realize logical scenarios of use and broad sub-classes of notification systems. To apply these concepts, we show how system sub-classes can be associated with various HCI resources, such as information and interaction design guidelines, example systems, prototyping techniques, and evaluation methods.

In developing a conceptual framework, our primary objective was to provide a taxonomic abstraction of user goals that preserves fundamental cognitive and perceptual/motor processes, yet allows easy and consist classification of nearly any notification system. With this, we gain advantages of reusable design knowledge and methods, but perhaps lose distinction of subtle cognitive processing demands. We started out by analyzing the drivers of success or failure for notification systems, recalling that these systems provide additional or secondary information to a user that is engaged in other tasks. We note that the difference between a successful or unsuccessful system ìhinges on accurately supporting attention allocation between tasks, while simultaneously enabling *utility* through access to additional informationî (McCrickard & Chewar, 2003), suggesting a cost-utility tradeoff. To describe the many possible sources of *utility*, we considered a wide variety of general tasks that notification information could support, and we feel the selection of three critical parameters (interruption, reaction, and comprehension) provides a reasonable approach.

For example, as users continue to attend to other concerns (e.g., editing a document or interacting with group members), they may desire notifications to monitor and understand how related or unrelated information states change over time, detect patterns and trends, slowly assimilate complex information, or gain awareness of the actions of collaborators. Goals such as these can be more generally described as *comprehension*ó in the sense that new information is processed and stored in memory. This differs from a general *reaction* goal, in which users want to receive a notification stimulus (e.g., a message that a new item has been added to a discussion forum) and immediately respond, with or without shifting their attention from the primary task or recalling the meaning of the information later. Examples of reaction include making quick decisions, modifying an approach to a primary task, proving a response, acknowledgement, or some other type of action. We also recognize that users may or may not value general forms of *interruption* as result of notifications. Interruptions can pace daily activities, prompt task transition to more critical tasks, allow receipt of urgent or timely information, or maintain synchronization with collaborators. Certainly, each of these three general goals can be a primary or exclusive goal of a notification, or other goals can cascade from an initial goal. For instance, a user may have a notification goal to be interrupted to receive important information (e.g., when a portion of a collaboratively authored document has been completed by a colleague), which then may lead to comprehending trends or causes and even a reaction goal (e.g. realizing colleague contributions or sending a quick acknowledgement signal). Likewise, another user may desire not to be interrupted, while still gaining an understanding of patterns or trends over time. Recognizing the user goals that a system should support is essentialó information presentation can drastically affect interruption, reaction, and comprehension levelsó often in a manner that sacrifices one for the other(s). Articulating the desired blend or transition of goal-states can help achieve the right tradeoffs.

To simplify a description of the user goals that a notification system supports, we recognize that the user desire for that goal (as well as a second measure of a systemís realization of each notification outcome) can be *high* (1), *low* (0), or any intermediate value. While we are currently investigating various objective methods of obtaining consistent descriptions, even a subjective judgment provides a useful articulation that can be understood by others and associated with development resources. Furthermore, this allows a compact annotation, which we call the *IRC rating* of a system or system state (I=interruption, R=reaction, C=comprehension). Therefore, if User A wants to maintain focus on a primary task (low interruption) yet send an acknowledgement message when a colleague submits a project resource (high reaction) without desiring any long-term understanding of the collaborator activity patterns (low comprehension), this design model has an IRC rating of 0/1/0. Alternately, User B may value notifications that are presented in a way that allows deep understanding of information over time and appreciation of patterns and trends, perhaps to learn about the work habits of colleagues (high comprehension), without causing primary task interruption or any type of near-term reaction (low interruption and reaction); this is reflected with a rating of 0/0/1. As a final example, User C may value occasional interruptions from a primary task in the form of friendly software agents or animations, perhaps by bringing humor or stress-relief to tedious tasks. This userís goals related to the comprehending or reacting to this information are low, so the IRC rating would be 1/0/0. The IRC framework, shown in Figure 1, allows IRC ratings to be expressed visually, depicting each critical parameter as an axis that ranges from low to high. The figure expresses each of the three example users and systems just described.

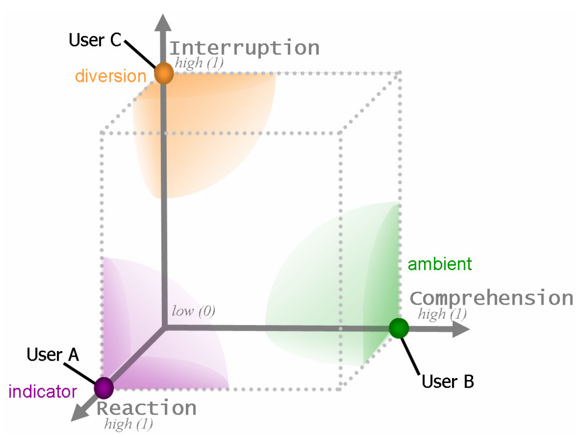

Figure 1: Framework reflecting the user goals for interruption, reaction, and comprehensionó critical parameters for system success. IRC-ratings are shown for the three example users and corresponding system sub-classes.

Representing combinations of user goals visually helps realize logical scenarios of use and broad sub-classes of notification systems. Systems with IRC-ratings close to the 0/1/0 corner (supporting User A) would be expected to provide a passive indication of an information state that could be quickly and easily interpreted and reacted to, so we refer to this corner as *indicators*. As we consider the intrinsic information design properties that would be necessary to support User Bís goals, it reminds us of *ambient* systems, which are designed to present meaningful information in a peripheral and calm manner. User C seeks a *diversion* from the primary task, nothing more than a brief and possibly enjoyable attention shift. We have applied a similar process to the other corners of the ìcubeî and identified general sub-classes and exemplar systems. Other sub-class names and regions appear later in Figure 2.

This framework can provide visualization of not only the design model (designerís interpretation of user goals) of a notification system, but also the *userís model* (the effect of the system from the userís perspective). A perfectly designed system would have two identical IRC-ratings, but user and situational differences as well as poor design choices can cause disparity. For example, although the designers of the notification system for User A may have been targeting low interruption and comprehension with high reaction (0/1/0), User A may perceive the system to be highly interruptive and may gain a great deal of unnecessary comprehension from it (although reacting as desired). Therefore, the userís model can be described as 1/1/1. Connecting the design and userís model plots reveals the reengineering approachó although the most effective approach may not always be a straight line (as depicted in Figure 2).

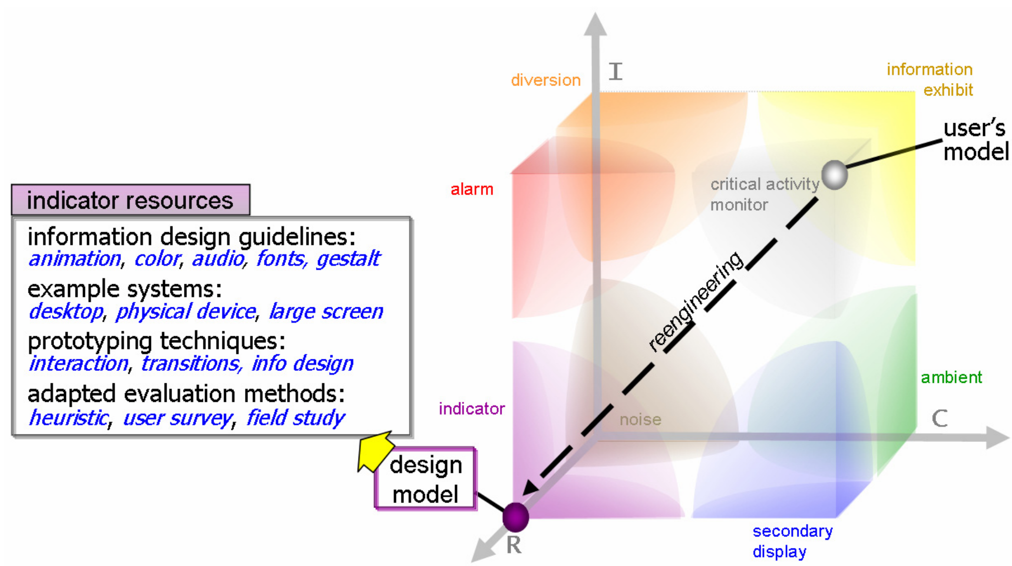

Figure 2: Design model and userís model plots of a notification system, revealing the reengineering approach necessary. Reengineering approach trajectories may be reusable, revealing common design problems. Design models can be associated with development resources that are specifically adapted for the combination of user goals.

The true utility of the IRC framework comes with the efficiencies provided through design model associations. Notification system sub-classes, such as ambient, alarm, or secondary displays, are compatible with user goals when they employ effective information presentation techniques (e.g. use of color to highlight or subdue, use of audio to attract attention, or use of animation to cycle information). Different forms of user interaction can affect IRC ratings as well. Guidelines for design elementsó laboriously extracted from human factors and experimental psychology literature or empirically developedó can be associated with regions within the framework. We have classified nearly one-hundred notification systems, all of which can serve as example systems, demonstrating varying effectiveness for different design models. Similarly, we have identified a range of prototyping and evaluation techniques that are being specifically adapted to portions of this design space. By articulating the design model for a new system, a developer can access this wealth of reusable knowledge resources, allowing significant development cost savings and increased systems usability. As the framework is used to develop notification systems, we can also collect data on reengineering approaches, relating specific design-user model disparity trajectories to redesign decisions and results. Over time, this can form reengineering patterns that can be analyzed for cost effectiveness in terms of impact on usability and user satisfaction, again providing reusable redesign options for similar design-userís model disparities.

To develop the framework more fully, we are exploring ways to represent systems that support a persistent rating, as well as task or information-driven rating changes. For example, a user may expect an ambient IRC-rating (0/0/.7) until an important information state is reached. This is expected to trigger a system change, presenting information that is suitable for an alarm IRC-rating (1/1/.3), which eventually transitions to secondary display behavior (.3/1/.7) and back to ambient. We are also considering how to represent systems that try to support simultaneous sets of user goals, such as quickly detecting and reacting to important new information during passive monitoring of overall patterns and trends. Of course, as an associative reference to potentially huge catalogues of design resources, the work can never be complete. However, we felt that the emerging idea of the framework could be particularly well suited for educating undergraduate and graduate students about the HCI processes and development concerns related to notification systems. To this end, we developed and executed three different versions of semester-long seminar curricula integrating the IRC framework.

Underlying Learning Program & Theory

In planning our semester-long seminar curricula, we felt that introducing the IRC framework to undergraduate and graduate computer science students could have two potential benefits: *1*ó it would help students appreciate attentional and multitasking concerns for notification systems design and facilitate access to useful resources, and *2*ó since the development of the framework is ongoing, it would inspire students to make contributions through testing concept applications, refining application procedures, and mapping the general critical parameters to domain specific parameters. While we discuss some of the outcomes of the IRC framework integration in the next section, this section introduces the approaches we took. We selected various activities based on our working knowledge of learning theory, so we describe these approaches in that contextó hopefully so that the general ideas can be useful to others. Since we have a longer-term interest in adapting the IRC framework to designers in industry, we also hope that the members of the ED-MEDIA community can deepen our insight about methods of effectively integrating the IRC framework into other instructional methods.

Our students are organized into three groups of mostly computer science students that meet in separate weekly sessions for about 1-1/2 hours. *Seminar A* consisted of ten undergraduate students, ranging from sophomore to graduating seniors with mid to high GPAs. Most had an introductory class on HCI, but little or no appreciable research experience. Only one of the ten was actively interested in graduate school, but for most the seminar satisfied an upper-level program requirement. *Seminar B* consisted of five students that are pursuing or planning to begin our 5-year BS/MS program, with the thesis-option. Each of these students (most are in the honors program or have very high GPAs) had previous HCI independent study research experience and expressed interest in developing a thesis related to notification systems. While some of these students were motivated by independent study grades, others were simply taking graduate research hours. *Seminar C* was open to any graduate student in computer science, as well as graduate students from our industrial and systems engineering department. While attendance for some sessions included up to 18 students, we attracted a core group of about 10. Most were first-year graduate students that had not yet selected an advisor or specific research interest, but a few others were activity engaged in notification systems-related graduate research.

To support the potential benefits of integrating the IRC framework into the three seminars listed above, we recognized two primary roles for the framework (corresponding with each benefit): *1*ó a schema that allows knowledge packets to be represented and used (Rumelhart, 1980) in conjunction with the popular HCI approach of scenario-based design (Rosson & Carroll, 2002), and *2*ó a context for situated cognition and learning (Lave & Wenger, 1991) that exemplifies a community of practice (Wenger 1998). We felt that the IRC framework could be effectively employed as schema if we included presentations and activity that focused on Rumelhartís processes of accretion, tuning, and restructuring and well as the suggestion of including schemata exercising in the form of problem-solving (1980). To develop a community of practice necessary for situated cognition and less-formal, opportunistic application of the IRC framework, we recognized the importance in perceptions of mutual engagement and accountability and a shared repertoire of practices, symbols, and artifacts (Wenger 1998). Throughout the design of our programs, we were especially sensitive to GagnÈís taxonomy of learning outcomes, particularly his conception of verbal information (providing an adequate, paraphrased description of schema instantiation), intellectual skills (developing a learning hierarchy involving discrimination of features, identification of classes, classifying examples, applying higher order rules), and cognitive strategies for visual expression (adapted from Driscoll, 2000). Supporting these outcomes requires consideration of critical learning conditions, such as creating meaningful but various context, using cues and distinctive features, and providing occasions for practice (GagnÈ & Driscoll, 1988). As a final concern, we hoped to allow a constructivist approach that promotes self-regulation of learning, where learners identify and pursue their own goals (Driscoll, 2000). Recognizing that there are enormous tradeoffs between these priorities and that our student groups each had slightly different needs, we selected three distinct approaches for integration of the IRC framework and general program design. The rationale for each is described below.

For Seminar A, we focused on supporting GagnÈ and Driscollís critical learning conditions to optimize creation of schema creation, accretion of notification systems development considerations, and tuning of previous HCI-related schema to fit the IRC framework. We addressed this by allocating much of the time in the first half of the semester to brief, highly illustrated presentations on aspects of the IRC framework and review of general HCI material, requiring completion of weekly exercises. These typically involved searching for concrete exemplar systems and resources as well as creating short, coherent arguments about necessary IRC framework modifications and limitations. As an example, we had each student develop specific steps for scoring the components of the IRC rating. This required deep consideration of the three cognitive processes (interruption, reaction, and comprehension), information presentation characteristics, and structured assessment methods. During the seminar, students explained their methods and practiced using each otherís methods to score systems and develop group consensus about ìthe method with the most potential.î We hoped that these exercises would build the community of practice, particularly since students would understand that their efforts were populating and improving the IRC framework. To enhance this feeling of mutual engagement, we brought in discussion of our IRC framework research activitiesó such as workshop and conference participation, running experiments, and presenting work to practitioners. We feared that this structured approach and the frequency of the exercises would detract from our constructivist ideals, particularly with respect to the quality of the semester research project. However, we thought that developing well-defined schema and inbound learning trajectories (Wenger, 1998) would be a beneficial tradeoff for this group.

We held an opposite conclusion for the Seminar B groupó since these students were already inbound or legitimate participators, we thought developing the community of practice was unneeded and instead, strongly favored the constructivist outcomes. At the seminar meetings, time was divided up among students to present their current progress, ask specific questions, and receive feedback from the group and supervisors. GagnÈís critical learning conditions were largely unfulfilled. To leverage situated cognition and reinforce the community of practice, we employed a cognitive apprenticeship model (Honebein et al., 1992) in which most students participated in more senior graduate student research objectives, receiving mentorship through weekly individual meetings, practicing actual research skills, but without full accountability risks. Although the students were expected to integrate the IRC framework into their research products, they were not provided any formal instruction for schema development, but were given access to papers (e.g. (McCrickard & Chewar, 2003)) and presentation resources used for Seminar A. Instead, we favored a focus on schema accretion and restructuringó particularly to extend our thinking about generic critical parameters to domain specific parameters, and attempted to create opportunistic and informal learning opportunities.

The overall objective was also slightly different for participants of Seminar C: while Seminar A involved a blend of structured learning and research production and Seminar B favored research production, Seminar C emphasized learning. As new graduate students uncertain about research interests, we thought to best serve in supporting development of schema and conversion of peripheral participation to an inbound trajectory within our existing community of practice. Like Seminar A, we decided that formal, highly illustrated presentations of IRC framework aspects would be best, but we favored longer presentations that also integrated scholarly HCI papers. While out of class activities would not be practical, in class activitiesó especially involving groupwork seemed feasible. We also saw an opportunity to introduce these students to the community of practice developed in Seminar C by exposing them to the activities and exercise products of the undergraduates. For example, after the Seminar A exercise (described above) that involved development of IRC rating procedures, students in Seminar C would be able to critique the groupís product and suggest even further refinement. Our hope was that this practice would develop inclinations toward mentorship, necessary for the sustainment of the cognitive apprenticeship model. Other activities involved active debate, pushing students to increase engagement by developing opinions and arguments. We also thought that allowing time to discuss personal research ambitions would strengthen feelings of inclusion and reinforce individual identity. Wanting to retain some constructivist flavor, we employed the approach of welcoming sporadic, ìcome when you wantî attendance. Although we placed a much greater emphasis on schema development and transfer, we felt that many of GagnÈís critical learning conditions were impractical for this group.

All three of the seminars ran for a full 16-week semester. To assess how well we integrated the IRC framework in seminar programs to fit the diverse learning needs of our students, we relied on a combination of evidence, discussed in the next section.

Educational Outcomes

How did the IRC framework impact undergraduate and graduate student understanding of notification system design challenges? We were quite pleased with the general outcomes of all three seminars, particularly with the energy surrounding the IRC framework. Each program version exhibited strong and weak pointsó often contradictory to our expectations. To frame a discussion of the outcomes, we present four primary sources of evidence: observed student attitudes and action, quality of research products, quantitative survey feedback, and expressed intent for further research. We discuss each in turn, highlighting differences between the three groups, inconsistencies with theory-based expectations (e.g. situated cognition, schema, constructivist), and outcomes that may be attributable to the integration of the IRC framework.

As we observed student attitudes and activities through their participation patterns and contributions, we were hoping to see development or maintenance of a community of actively engaged notification system researchers. All three seminars succeeded in this regard, although Seminar A exhibited the most change and seemed to achieve the deepest sense of community in the end. Although initial class participation was low, as IRC-related schema was developed and exercised, participation levels grew stronger. While the same trend was observed for the graduate students in Seminar C, the change was smaller. The attitudes of these students can be characterized as deeply reflective of the material, while the undergraduates were eager to question, challenge, and extend the IRC framework. As schema development progressed, we sensed a feeling of pride within the Seminar A studentsó the seminology and higher order rules associated with the IRC framework may have been perceived as acquired tools that were esoteric but necessary for functioning within and contributing to the shared repertoire of the community. Secondary effects were encouraging: students voluntarily participated in each otherís research, helping to run experiments or act as participants, and everyone exhibited a strong sense of accountability for their research activities. Toward the end of the semester, one particularly poignant example of the effect of the IRC framework was level of critical analysis and problem introspection demonstrated in a debate about usability evaluation methods most suitable for each corner of the IRC cube. However, observed attitudes and activities within Seminar B show a weakness in the IRC framework integration approach. While some students sought out knowledge about the framework and ways to tie it to discussions about ongoing research, others were not self-motivated or interested in obtaining this knowledge. At times, this seemed to cause a fracture within the group, temporarily creating a condition of detachment but perhaps diverting some from their inbound trajectories over the long-term. This same effect was observed for Seminar C students that were not regular attendees.

We expected the quality of research products from Seminar B to surpass those from Seminar A, largely as a function of academic record and motivation for graduate study. However, we observed more indication of the oppositeó students from Seminar A produced outstanding research products, despite the effort involved with the weekly exercises, the lack of accountability associated with semester-long weekly progress reports, and the lower amount of individual feedback and mentorship. Although we thought the conditions created for Seminar B would promote self-regulation of learning, we found procrastination and lateness to be a much larger problem with this group than Seminar A. Students in Seminar A selected their own research topics, but all were able to effectively use or extend the IRC framework. Topics included an adaptation of analytical evaluation heuristics for notification systems, development and evaluation of collaborative notification systems with various design model IRC ratings, survey of notification system design trends in homes throughout the past 20 years (using the framework to show changes in design and userís model ratings), lab-based comparison of audio and visual cues to understand userís model notification characteristics, and analysis of and prototypical design space population for mail notification systems. All Seminar A students expressed interest in submitting their efforts to conferences, and nearly all of the non-graduating students made plans to continue notification systems research during the next semesteró suggesting unexpected constructivist outcomes.

Although we elicited feedback in the form of a survey to better understand how specific activities and overall outcomes were viewed by the students, we are unable to tie the results to specific IRC framework integration choices or group differences. However, more than half of the responding students across all three seminars selected some direct facet of the IRC framework as an aspect of the seminar (presented among other choices) that contributed to their learning in five of six stated objectives (supported objectives were: understanding notifications systems, describing what notification systems should be providing users, considering interface design options and challenges, and comparing methods and options for usability evaluation; the unsupported objective was implementing early design ideas). The only surprising results from the survey were a consensus among Seminar A students that time spent during the seminar and outside the seminar on exercises should not be changed (not reduced), that time spent on the presentations (reviewing HCI material and introducing the IRC framework) should be decreased, and that more time should be spent during the semester on projects. However, students in all programs at least agreed that their understanding of HCI was extended by the seminars and that they found it to be an effective mix of theory and practical application for designing and evaluating notification systems.

Conclusion & Future Work

Overall, we can conclude that the IRC framework was most effective when presented in detail to the full group. After introducing the framework to students in those conditions, follow-on exercises, discussions, and projects showed rich and effective employment of intellectual skill outcomes. The framework certainly initiated discussion and provided a structure for critical analysis. We feel that we were able to use it effectively to create and broaden a community of practice necessary for situated learning and meaningful action. Based on this, we believe that the IRC framework has vast potential for lending efficiency to the design and evaluation of notification systems as well as supporting HCI education for these systems. We want to find new and better ways to integrate the framework into instructional activities, founded on solid learning theory. We are also interested in creating and evaluating notification system applications that support education in a dual-task context, and we are looking for opportunities to collaborate with others on such projects. We perceive this future work to be an incubator for continuous improvement of our conceptual framework and the associated catalog of reusable design knowledge and resources.

Acknowledgements

We thank all of the students that participated in the seminars for their hard work. Jacob Somervell and Ali Ndiwalana also provided assistance in developing the IRC framework.

References

Carroll, J. M., Neale, D. C., Isenhour, P. L., Rosson, M. B. & McCrickard, D. S. (In press). Notification and awareness: Synchronizing task-oriented activity. *International Journal of Human-Computer Studies*.

Driscoll, M. P. (2000). *Psychology of Learning for Instruction* (2nd ed.). Allyn & Bacon, Boston.

Erickson, T., Halverson, C., Kellogg, W. A., Laff, M., & Wolf., T. (2002). Social translucence: Designing social infrastructures that make collective activity visible. *Communications of the ACM 45*, 4 (April), pp. 40-44.

GagnÈ, R. M. & Driscoll, M. P. (1988). *Essential Learning for Instruction* (2nd ed.). Allyn & Bacon, Boston.

Honebein, P. C., Chen, P. & Brescia, W. (1992). Hypermedia and sociology: A simulation for developing research skills. *Liberal Arts Computing 1* (pp. 9-15).

Lave, J. & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, New York.

McCrickard, D. S. (1999). Maintaining information awareness with Irwin. In *Proceedings of the World Conference on Educational Multimedia/Hypermedia and Educational Telecommunications (ED-MEDIA '99)*.

McCrickard, D. S. & Chewar, C. M. (2003). Attuning notification design to user goals and attention costs. *Communications of the ACM 46*, 3.

Rosson, M. B & Carroll, J. M. (2002). Usability Engineering: Scenario-Based Development of Human Computer Interaction, Morgan Kaufmann, New York.

Rumelhart, D. E. (1980). Schemata: The building blocks of cognition. In R. J. Spiro, B. C. Brace, & W. F. Brewer (Eds.), *Theoretical Issues in Reading Comprehension*. Erlbaum, Hillsdale, NJ.

Wenger, E. (1998). *Communities of Practice: Learning, Meaning, and Identity*. Cambridge University Press, New York.

Wickens, C. D. and Hollands, J. G. (2000). *Engineering Psychology and Human Performance*. 3rd ed. Prentice Hall, Upper Saddle River, NJ.