Towards Notifications for Mobile Response Teams

Peter Tarasewich
College of Computer and Information Science
Northeastern University
Boston, Massachusetts 02115 USA
tarase@ccs.neu.edu

Christopher S. Campbell
IBM Almaden Research Center
650 Harry Road, 4CP/E3
San Jose, California 95120 USA
ccampbel@almaden.ibm.com

Abstract
Mobile response teams, such as firefighters and combat personnel, encounter fast-paced, information-intensive situations with high rates of contextual shift. These situations require precise communication and coordination of actions between team members. This paper reviews our work to date concerning context modeling and the design and testing of pixel-based visual notification cues. We then present current research involving the design of mobile technology to support mobile response teams and their activities.

1.0 Introduction
Mobile notification technology has the potential to provide effective ways to coordinate team members, increase communication and awareness, and ultimately improve performance and productivity. This potential, however, will only be realized by understanding how to design notification systems considering human cognitive abilities and goals, the tasks performed, and the situation or context. With mobile technology, these constraints become even more severe as cognitive abilities are stretched to the limit, tasks are many and varied, and context rapidly shifts. Information overload can also be a serious problem as multitudes of dynamic sights, sounds, and other stimuli compete for a person's limited attention.

One way to reduce information overload is through the use of meta-information, which can require less effort to process and can result in fewer or less severe disruptions. If meta-information is deemed important, the person receiving it can make a decision whether or not to seek additional details. For example, a mobile worker may not need or want the entire contents of a message every time one becomes available. It may be too distracting (and perhaps too dangerous) to the worker's primary tasks. However, they may wish to receive a notification that a message is available, along with its importance and source. That way, the worker can make a decision based on their current situation whether to stop their primary task to access the message contents.

Pixel-based visual notification cues use one or more individual lights to indicate the status or availability of information that is of interest to a particular user in a ubiquitous setting. When notification cues are used to convey meta-information, questions arise such as (a) what form should they take, (b) when should they be sent, and (c) how appropriate they are given the recipient's current situation and activities. The design of mobile notification cues is therefore a complex problem, requiring the selection of appropriate delivery channels based on continuously changing contexts and dynamic information needs [7].

The practical benefits of pixel-based visual notification cues are many. These cues (e.g., in the form of LEDs) can be embedded almost anywhere. An output screen is not needed for such cues, allowing the display of information on very small, or *ultra-mobile*, devices. In addition, the transmission of such cues requires relatively little bandwidth, and their display requires minimal power. Unlike text messages or icons, pixel-based cues can be personalized such that only the user knows what they mean, even if they are publicly displayed. For example, three yellow lights on a ring, even when noticed by other people nearby, could convey a message only understood by the wearer. Thus, high levels of security and privacy can be ensured with pixel-based visual notification cues. These cues also allow for language-independent communication, aiding in the quest for universal usability.

Much of our current research focuses on using mobile technology to increase productivity by improving the coordination and communication of *mobile response teams* (MRT) and their activities. Notification cues provide awareness of incoming messages and changes in status so that teams can organize themselves and respond to a changing environment. MRT often encounter fast-paced, information-intensive situations with high rates of contextual shift, requiring precise coordination of actions between their members. Examples of MRT include firefighters responding to a natural disaster and military personnel conducting a combat mission on foreign soil.

This paper provides background on our previous research in the areas of context and visual notification cue design and implementation. The paper then discusses the start of our work in designing notification cues (visual or otherwise) that can support the activities of mobile response teams.

2.0 Background

We present a brief review of recent work by ourselves and others in the areas of context modeling, notification cue design, and context-awareness.

2.1 Context

Before mobile devices, most users worked with computers in contexts that remained relatively stable. Contextual concerns could be taken into account during the design process and changed little, if at all, after system completion. With the advent of mobile devices, context-appropriate design has become more difficult, for reasons such as:

a) The context may change quickly over time
b) The user's priorities and tasks may likewise change unpredictably and rapidly
c) The context may be unfamiliar to the user

We have defined a comprehensive model of context (built on the strengths of previous models ([e.g., 1, 11, 14]) using three broad categories: environment, participants, and activities [15, 16] (see Figure 1). The "environment" category is concerned with the characteristics of the physical environment and objects within that environment. "Participants" includes the status of the user(s) and other people in the environment. This also encompasses personal properties such as age, gender, preferences, skills, and physical capabilities [5]. The "activities" category covers user, participant, and environmental activities. Representative characteristics are listed in Table 1. In addition, the model reflects interactions or relationships that may exist between participants, activities, and environment. This includes the co-location of people, group dynamics, and social situations. The categories are considered over a timeline of past, present, and future. This allows for a context history, which can be used for comparison to the current context, or for predicting future context.

Each of these aspects of context affects decisions about notification cue design and implementation in different ways, although there are relationships between the different pieces as well. Location can change rapidly and will influence what type of notification cue is most effective (e.g., visual cue in a noisy environment), appropriate (e.g., vibration in a library), or even safest (e.g., voice cue for a driver). Properties of a location can affect presentation of a notification cue (e.g., brightness on a sunny versus cloudy day). Activities and the co-location of people at a specific location (along with the relationships between the people) will determine the effectiveness and appropriateness of notification cues (e.g., an auditory cue for someone talking with friends outdoors and a tactile cue when speaking with colleagues in an office).

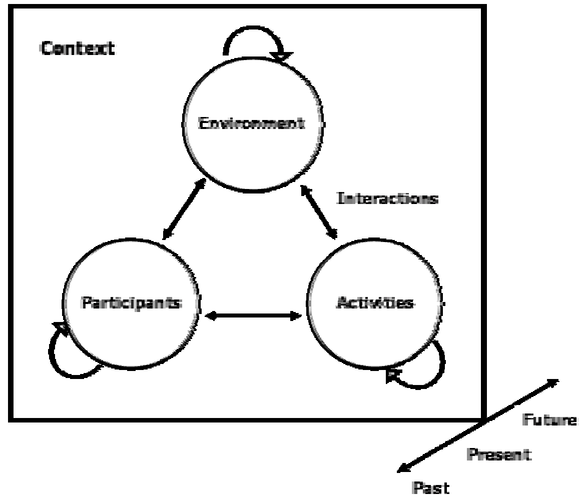

Figure 1 – Context Model

Category	Representative Characteristics
Environment	Location, Orientation (of objects) Physical properties Brightness and noise levels Availability, quality (of devices and communications)
Participants	Location, Orientation Personal properties (e.g., age, gender, education, preferences) Mental state Physical health Expectations
Activities	Tasks and goals (of participants) Events in the environment (e.g., weather)
Interactions	Co-location Group dynamics Social situations Participant/environment relationships (e.g., worker/workplace) Season, time-of-day, day-of-the-week

Table 1 – Characteristics for Context Categories

2.2 Notification Cues

While conveying relatively little detail, notification cues provide meta-information, or information about information, to their intended recipients. This meta-information may be as simple as a signal indicating an incoming cell phone call, or as complex as a signal providing the priority, the sender, and a summary of a new email message. On ultra-small devices (e.g., watches or jewelry), it may be feasible to convey certain information only in the form of notification cues. To be successful, a cue must 1) get the intended recipient's attention, 2) clearly convey its information to the recipient, 3) minimize the disruption of the recipient's current task(s), and 4) minimize the disruption of other people near the recipient.

2.2.1 Getting the intended recipient's attention

Notification cues must safely compete for a user's attention in a world full of distractions. While sensory overload is one concern for notification design, another is subtlety. There are also risks of cues being ignored because of inappropriate or overused designs. Critical notifications that fail to be delivered can lose their value completely.

Context affects the appropriateness of a notification cue and the probability of successfully getting a recipient's attention. For example, if the recipient is on a busy city street with two small children during a rainstorm trying to find a doctor's office, a visual or even an audio notification may not work well. In this situation, a tactile notification might work best (e.g., a vibrating ring or watch). Schmidt et al. [13] studied context-appropriate notifications using a cell phone that changed its ring type (e.g., vibrate, quiet ring) based on its perceived context. Sensor data was interpreted for information about the phone's current situation, and its ring was adjusted accordingly.

2.2.2 Clearly conveying information to the recipient

Given the many forms (i.e., visual, audio, tactile) that notification cues can take, it is necessary to study which forms can be comprehended most easily and consistently, and how context affects comprehension. The amount of information that can be successfully carried by different notification cues under varying circumstances must also be determined.

We conducted a study that measured the performance/size tradeoff of visual displays that ranged in size from two lights to nine lights, and used display characteristics, such as color and blinking, in various combinations [18]. Results showed a reliable tradeoff between performance (response time and accuracy) and display size (number of lights). However, even the full set of twenty-seven messages used in the study could be conveyed with high recognition accuracy using only three lights by mapping the messages into color and position. We concluded that micro-sized devices could be designed to convey critical information and provide effective notifications. However, this study did not explore 1) how learning affected the comprehension and use of the displays, and 2) how much information could be effectively conveyed using a given display size.

A follow-up study [2] investigated these two open issues through an experiment that measured user learning and comprehension of increasing amounts of information on a three light visual display. Each light displayed one of three colors at one of two intensity (brightness) levels. Users were required to learn five sets of messages of increasing size and complexity using the display. Results showed that micro-displays could transmit detailed, information-rich messages up to 6.75 bits in size with minimal training, that is, few trials and short time frames.

A third study [17] performed field-testing of notification cues on a mobile handheld device. Each cue consisted of three multicolored lights preceded by a tactile signal (vibration). After being customized by a user, the cues were sent periodically to the device over a wireless network as the user went about their normal activities. User personalization seemed to enhance learning and usefulness of the cues, while the multimodal design aided arrival awareness.

Other researchers have looked at using colors and patterns of pixels to display information, albeit on a much larger scale. Prante et al. [9] discuss the development of Hello.Wall, a wall-sized ambient display consisting of a grid of 124 cells (LED clusters). The wall displays information as colored light patterns based on the proximity of different users to the wall.

Tradeoffs may exist between the amount of information that a notification cue can successfully convey and the context of the intended recipient. Certain contexts may necessitate or favor multi-level cues, where an initial cue conveys a single dimension of information (e.g., priority), and the user can call up additional details (e.g., source) before responding further. While cues can provide multiple levels of information at once (e.g., priority, source, and content), in certain contexts (e.g., high-stress environments) it may be better to convey only part of that information initially.

2.2.3 Minimizing disruption of current task(s)

Notification design has to take into account the intended recipient's current task, how that task is being performed, and any additional cognitive limitations that are currently placed on the intended recipient. For example, if a visual task is being performed, then an audio notification cue (which is dissimilar to the current task) might be appropriate.

2.2.4 Minimizing the disruption of other people

While it is important to convey a notification to its intended recipient, it is also important to minimize the disruption of other people in the vicinity of the recipient. Notifying a person who is in a social meeting with friends might be handled differently than one who is in a business meeting with colleagues, in a sales presentation with potential customers, or in a strategic planning session with senior executives.

It is also important under certain circumstances to protect the privacy and security of information sent to people through notification cues. With any type of notification cue, it is possible to allow some personalization of the cue through customization. With pixel-based visual cues, a number of different characteristics might be user selected---e.g., color, intensity, position, and blinking. The user could also select the mapping between cues and messages. In this manner, only the user would know the message even if many others see the display. Privacy and security advantages of using colored pixel-type displays are also discussed in [9]. One potential problem with this approach is that the user may select cues and mappings that are difficult to remember.

2.3 Context Adaptation

A truly adaptive mobile system would take into account relevant changes in the user's environment on a real-time basis and modify notification cues as appropriate. While seemingly straightforward in theory, in reality such systems are difficult to implement for a number of reasons. First, determining which context factors are relevant to a user and application at any time is difficult. Second, technologies that adequately sense and process changes in the environment are still under development. Third, one needs to determine if (or how) context information should be used to change the form of a notification cue before it is delivered.

Devices that derive input indirectly from the user might improve usability in a dynamic environment. Schmidt [12] discussed a vision of mobile computing where devices can "see, hear, and feel." Devices act according to the situational context in which they are used, and might receive input from their surroundings rather than from the user.

Adaptable systems, however, must still account for the limitations and needs of their human counterparts. Schilit, Adams, and Want [11] noted the potential for distraction when the context is changing rapidly, or the impracticality of adapting to every change. Confusion may also result from incorrect readings of context conditions. Overall, a set of strategies is needed for building context-aware applications. In this regard, Chevrest et al. [3] put forth the following:

- Do not inappropriately over-determine the interaction
- Adhere to the principle of least astonishment and the need for predictability
- Use dependable sensing technology
- Allow users to override any adaptation strategies

3.0 Mobile Response Teams

Through a review of existing case studies and research literature, we have started to investigate the information management and coordination issues that face three types of mobile response teams: firefighters [e.g., 6], emergency medical personnel [e.g., 4, 8], and combat personnel [e.g., 10]. There are some common attributes of the situations and activities for the three MRT described above. Team members:

- Must be able to communicate with each other
- Need to know each others' locations
- Must have hands and eyes relatively free to perform their primary tasks
- Must remain as mobile as possible
- Need to monitor changes in the outside environment
- Need to monitor changes in team or activity status
- Work in potentially harsh and dynamic environments
- Are often burdened with heavy equipment and/or supplies
- Are under large amounts of stress, and heavy cognitive loads
- Are watching out for their own lives, and the lives of others

Other factors that can affect some mobile response teams include:

- Communications and information must remain private and/or secure
- Noise can be a potential problem (either too much, or the need for silence)

Taking these factors into account, we can generate general guidelines for the design of notification cues to meet the needs of these types of mobile response teams. Note that these guidelines assume the availability of a reliable communications network. Notification cue display designs for mobile response teams should be:

1. **Hands-free.** Information should be readily available without the need to use fingers or hands to retrieve it.
2. **Light.** Mobile response team members are already overburdened with equipment and other things to carry.
3. **Usable under extreme conditions.** If visual cues are used, they need to be readable in darkness and sunlight. If audio cues are used, they need to be loud enough to be heard over any noise present in the environment.

4. **Simple.** Team members must be able to quickly and accurately comprehend cues that they receive.

In addition to this, notification cue displays will need to be designed for specific types of mobile response teams. For example, firefighters might benefit most from a simple visual display for certain information (e.g., evacuation orders) because of the noise from a roaring fire. Combat personnel on secret missions might also benefit from visual displays, but for the reasons that any noise might be dangerous to the team member and to the success of the mission.

4.0 Supporting Mobile Response Teams

The research described above provides the foundation for our current and planned studies. Some of these involve answering fundamental questions about the design of pixel-based visual notification cues for the mobile environment:

1. Which cue designs work best in different environments?
2. Should pixel-based visual notification cues be augmented with additional cues (e.g., tactile or auditory) under certain circumstances?
3. How can we model, create, and evaluate pixel-based visual notification cues that exhibit automatic real-time adaptation to a dynamically changing context?

One of our current projects looks more closely at how learning and comprehension of cues are affected by customization of the cues by their user. We are directly comparing performance with predefined cues against those in which the user chooses the cue designs, mappings, and/or messages. We are also expanding our studies to test various cues under more realistic conditions (e.g., for decision making purposes). This will begin to help us understand how well cues work in various activities, and how much the recipient's current task or tasks are disrupted (and the effect of this disruption).

The potential privacy and security aspects of pixel-based visual cues are one of their most intriguing and potentially useful features. Pixel-based cues can be personalized such that only the intended recipient knows what they mean, even if they are publicly displayed. While this can also be done to an extent with concepts such as custom cell phone rings for unique callers, we hypothesize that users can more readily comprehend information from a pixel-based cue than from an auditory (or tactile) cue. This is something we continue to investigate.

Our past studies have used simulation to test basic cue design and comprehension concepts. Future studies will use more realistic form factors (e.g., LEDs). We are also constructing more realistic prototypes, and planning longitudinal studies in which the user will interact with a given prototype over a period of days or weeks in order to better assess use of such cues in the real world.

Pixel-based visual notification cues (or visual cues in general) may not be the best choice under all circumstances. Our initial prototype study [17] showed that adding a tactile dimension to the visual notification aided awareness. Another possibility is preceding the visual cue with a sound. Sounds or vibrations might also be used in conjunction with the visual cues to transmit information (e.g., pitch indicates the priority of the message. Use of such "multimodal" cues may be advantageous or even necessary in many circumstances, and are being investigated further.

We are planning to develop context-sensitive devices to perform field studies to see how well notification cues can be adjusted based on the assessment of context (e.g., where the user is, who the user is with, what the user is doing). This will require 1) an assessment of which context characteristics are important to notification cue delivery, and 2) the development of algorithms to process context information and to determine the appropriate format for notification cues.

Parallel to this work on pixel-based cues and their appropriateness for different environments, we plan to perform laboratory and field experiments to test the usefulness of different cues for mobile response teams under simulated conditions, and eventually move toward real-world field tests.

Our work with pixel-based notification cues has shown them to be a very promising form of communication for ubiquitous and mobile environments. We also view notification cues (in whatever form) that adapt to a user's context as an equally important research issue. Mobile response teams, and potentially all mobile users, stand to greatly benefit from research in both of these areas.

References

1. Abowd, G. D., and E. D. Mynatt (2000). "Charting Past, Present, and Future Research in Ubiquitous Computing," *ACM Transactions on Computer-Human Interaction*, 7(1), 29-58.
2. Campbell, C., and P. Tarasewich (2004). "Communicating with Three Lights," *Extended Abstracts of CHI 2004.*
3. Chevrest, K., N. Davies, K. Mitchell, and C. Efstratiou (2001). "Using Context as a Crystal Ball: Rewards and Pitfalls," *Personal and Ubiquitous Computing*, 5(1), 8-11.

4. Holzman, T. G. (1999). "Computer-Human Interface Solutions for Emergency Medical Care," *Interfaces*, May/June, 13-24.
5. Jameson, A. (2001). "Modeling both the Context and the User," *Personal and Ubiquitous Computing*, 5(1), 29-32.
6. Jiang, X., N. Y. Chen, J. I. Hong, K. Wang, L. Takayama, J. A. Landay (2004). "Siren: Context-Aware Computing for Firefighting," *Proceedings of the Second International Conference on Ubiquitous Computing (Pervasive 2004)*, forthcoming.
7. McCrickard, D. S., and C. M. Chewar (2003). "User Goals and Attention Costs," *Communications of the ACM*, 46(3), 67-72.
8. Munoz, M. A., M. Rodriguez, J. Favela, A. I. Martinez-Garcia, and V. M. Gonzalez. (2003). "Context-aware mobile Communications in Hospitals", *IEEE Computer*, Sept. 2003, 38-46.
9. Prante, T., C. Röcker, N. Streitz, R. Stenzel, C. Magerkurth, D. van Alphen, and D. Plewe (2003). "Hello.Wall – Beyond Ambient Displays," *Adjunct Proceedings of UbiComp 2003.*
10. Sass, P. (1999). "Communications Networks for the Force XXI Digitized Battlefield," *Mobile Networks and Applications*, 4, 139-155.
11. Schilit, B., N. Adams, and R. Want (1994). "Context-Aware Computing Applications," *Proceedings of the Workshop on Mobile Computing Systems and Applications*, 85-90.
12. Schmidt, A. (2000). "Implicit Human Computer Interaction Through Context," *Personal Technologies*, 4(2&3), 191-199.
13. Schmidt, A., K. A. Aidoo, A. Takaluoma, U. Tuomela, K. Van Laerhoven, and W. Van de Velde (1999). "Advanced Interaction in Context," In *Proceedings HUC '99,* 89-100.
14. Schmidt, A., M. Beigl, and H.-W. Gellersen (1999). "There is More to Context than Location," *Computers & Graphics*, 23(6), 893-901.
15. Tarasewich, P. (2003). "Designing Mobile Commerce Applications," *Communications of the ACM*, 46(12), 57-60.
16. Tarasewich, P. (2003). "Towards a Comprehensive Model of Context for Mobile and Wireless Computing," *Proceedings of the AMCIS 2003 Conference*, 114-124.
17. Tarasewich, P., T. Bhimdi, and M. Dideles (2004). "Testing Visual Notification Cues on a Mobile Device," *Extended Abstracts of CHI 2004.*
18. Tarasewich, P., C. Campbell, T. Xia, and M. Dideles (2003). "Evaluation of Visual Notification Cues for Ubiquitous Computing," *Proceedings of UbiComp 2003*, 349-366.