

Change Blind Information Display for Ubiquitous Computing Environments

Stephen S. Intille

MIT Home of the Future Consortium
1 Cambridge Center, 4FL
Cambridge, MA 02142 USA
`intille@mit.edu`

Abstract. Occupants of future computing environments with ubiquitous display devices may feel inundated with changing digital information. One solution is to create a reasoning module that accepts requests to display information from multiple applications and controls how the information is presented to minimize visual disruptions to users. Such a system might use information about what activity is occurring in the space to exploit a powerful phenomenon of the human visual system: change blindness.

1 Introduction

One common vision of the future of computing is that human-computer interaction will transition from the desktop to the environment. Ubiquitous sensor systems will continuously observe human activity and infer context. That context will be used to present information to people at precisely the time and place that it is needed. The hope is that these *context-aware* systems will increase human-computer information flow without contributing to feelings of information overload [1].

“Calm” environments are those that do not trigger information overload [35]. For many people, the environment that fosters the most intense feelings of calmness, security, and relaxation is the home. Homes are often comfortable, friendly places that provide a respite from the unpredictable onslaught of information one can experience in other environments. Can context-aware technologies be brought into the home without destroying the home’s calm aesthetic?

Two user interface design strategies have been proposed for minimizing information overload. The first strategy is to use context to infer user intent and actively present information such as context-sensitive help (e.g. [10]). The second strategy is to shift information from the user’s focus of attention to the user’s visual and auditory periphery (e.g. [35]). Although both of these techniques might reduce feelings of information overload if they are executed well, they do not eliminate the underlying cause of a disruption of calmness. The *detection of change* itself may destroy the sense of calmness in a space, regardless of whether these changes are in the focus of attention or in the periphery. This paper outlines a

strategy for using context-aware computing to minimize perceived information change. The proposed method exploits fundamental properties of the human visual and memory system to minimize involuntary perception of user-interface-induced distractions as information is presented and updated. This technique may make it possible for context-aware sensing systems to provide information to people in ubiquitous spaces but to do so in a way that preserves the feeling of calmness associated with non-digital, everyday environments such as the home.

2 The problem

Most environments today do not have a sufficiently dense network of output devices to create strong feelings of information overload. Soon, however, it will be possible to create affordable office and home environments that have low-cost tiled wall displays (e.g. [12]), electronic ink displays [18], head-mounted displays, and/or augmented reality projection systems (e.g. [23]).

The ability to ubiquitously present information creates a design challenge. With few exceptions, present-day media – television shows, commercials, movies, web sites, software, and electronic signs – have migrated towards the use of more intense visual and auditory stimuli in shorter amounts of time. Two recent examples of this are CNN’s continuously scrolling news highlights and the use of graphics and animation in sporting broadcasts to emulate video games. Such trends suggest that stimulus-inducing information clutter in the home will only intensify with the introduction of more versatile ubiquitous display technology.

2.1 A present-day scenario

For instance, imagine that you are in a meeting with 10 other busy people, all of whom have cell phones, PDA devices, and notebook computers. Would you prefer that the other people in the room have their devices set for silent notification of incoming messages (e.g. via vibration) or audible notification (e.g. via unique rings)? Most people would prefer others to use silent notification because it is less disruptive.

The least disruptive meeting scenario occurs when a person receives a silent notification, makes no sudden movements, and calmly gets up and leaves the room at a time when that person is not actively engaged in the ongoing conversation. Nothing attracts (and thereby distracts) the attention of others in the space. There is a gradual, calm transition from the state before a message is received to the state after the call is answered without triggering innate human alerting responses of others in the room.

On the other hand, if a device suddenly rings, the human auditory and visual system ensures that everyone in the room will respond to the novel stimuli with a mild startle or orienting reflex [26]. Given enough such interruptions, people will habituate to the stimulus and the reflex will abate. Unfortunately, as these types of personalized messages become more common, individuals are personalizing their messages so that each has a distinct stimulus. Habituation with so

many signals becomes more difficult. Somewhat counterintuitively, greater personalization and targeting of information is actually increasing the disruptive impact of the information for others.

For receiving their *own* messages, people generally prefer audible notification because they are less prone to miss a message than with silent notification.¹ Although people are conscious of the growing social bias against disruptive, audio notification, the desire to receive an important incoming message is stronger than the social pressure to prevent distracting messages from reaching others.

2.2 A future (worst-case) scenario

The danger of constructing environments where information can be displayed everywhere is that, too often, information *will* be displayed everywhere. Figure 1a shows an architect’s rendering of a ubiquitous computing lab with a digital table, counter, and wall surfaces illustrating how it might look when all the displays are on. Much of the information could be dynamic, attracting the user’s attention. Figure 1b shows a real-world environment. New York City’s Times Square is the antithesis of a calm computing environment. Huge digital walls display television shows and ads, scrolling text, short animations, and images. Although invigorating for a short period of time, most people would not want their homes to continuously provoke a heightened stimulus response.




Fig. 1. (a) A simulated worst-case ubiquitous computing scenario where every available display is filled with information, much of it dynamic. (b) A real-world example in an urban space.

The trend towards attention-grabbing information clutter is a byproduct of the fact that many interfaces and information environments are not created based on a single, coherent vision of how information should be best conveyed

¹ Vibration mechanisms are only effective when mobile computing devices are pressed against the body. They also quickly drain batteries and are sometimes more startling for the intended recipient than auditory cues.

to people. Instead, the user experience emerges from *competing interests*. The amount of time a person has to absorb new information remains relatively constant, so competing interests escalate the intensity of visual attractors to attract limited user attention. Corporate websites are often cluttered for the same reason – competing interests in the company insist upon staking out a position from which to catch the user’s attention, despite scarce screen real estate.

Homes of the future will not be free of competing interests. Vying for the home occupant’s attention will often be in the best interest of the creators of devices that use or control visual and auditory displays. Future homes may not look like Times Square, but existing trends suggest that without additional constraint, dynamic and stimulus-inducing information will tend to creep into the environments, filling all available display space.

The question considered here is, what can be done about this?

3 Related work

Weiser and Brown identified the trend towards information overload and observed that “designs that encalm and inform meet two human needs not usually met together. Information technology is more often the enemy of calm” [35]. They go on to advocate the development of *calm technology* that “engages both the center and the periphery of our attention” and allows information to move “back and forth between the two.” Software that facilitates such transitions has been called the “sweet spot” of ubiquitous computer design [6]. The strategy is to use technology to create interpretive clues in the periphery that help people quickly make decisions, facilitate communication [5], and provide “locatedness” [35].

An architectural example used by Weiser and Brown to motivate the adoption of calm technology is the interior office window [35]. The claim is that the window helps the office dweller maintain awareness of office activity in a calm fashion that avoids direct messages that might create a feeling of information overload. Computer technologies have since been developed that use non-traditional output modalities to convey information in the user’s periphery. These “ambient displays” [36] range from a dangling string that sways with levels of Internet traffic [13] to bubble fountains that convey information [8] to light patterns and pinwheels that map to office activity outside a cubicle [36]. A challenge is to find natural mappings between information and the environment [21].

Ambient display research presupposes that changing information in the user’s periphery preserves a sense of calm better than alerting the user of changes directly and that peripheral display will achieve the goal of putting “us at home, in a familiar place” [35].² The user, however, must still cognitively process these changes. For instance, although interior windows can help people maintain a

² Researchers advocating the use of peripheral displays to promote calm implicitly assume this strategy will promote calm environments. The author has not found direct psychophysical evidence for the conditions required to maintain a “calm” environment. This is an issue needing future work.

sense of connectedness with nearby activity, many people must still close the blinds and shut the door to minimize peripheral cues in order to concentrate. Cubicle dwellers express frustrations at the peripheral cues that they must endure.

Most proposed uses of peripheral information are for “non essential” information given the current task [17] that can often be presented in an abstract form. Peripheral display has been used to supplement desktop displays with the goal of improving the multi-tasking abilities of users [16]. Not surprisingly, peripheral “ticker” displays that use continuous scrolling have been found to be more distracting for users than tickers using discrete scrolling. Each ticker change attracts the users attention away from the primary task [17].

4 Creating calm by minimizing perceived change

An alternative (or complimentary) strategy to create calm environments is to minimize *all* attention grabbing cues presented to the user – in the focus of attention or in the periphery. It is possible that any *changing information*, regardless of what form it arrives in, may erode the perception of calm.

As indirect evidence of the relationship between a feeling of calm and the minimization of detectable change, perform a thought experiment. Imagine that you are in Times Square or the mock environment in Figure 1a. In the first condition, wherever you look information is static but in your periphery you sometimes detect information changing. In a second condition, wherever you look all information is static but yet somehow all the information is up to date. Most people think the latter condition would be more conducive to creating or maintaining a calm environment.

Our visual systems have evolved to keenly detect sudden changes in the visual field. Visual stimuli changes trigger the eyes to rapidly move, or saccade, in order to attend to the “motion transient” [29]. This makes sense from an evolutionary perspective: normally change causes motion [14]. The benefit of digital information is that it can be instantly and automatically updated, but each visual change can create a new motion transient. Our biology compels us to attend to these motion transients.

Each person in a future computing environment may want to have information displayed nearby that (1) stays current, (2) is sometimes augmented on the relevant objects as the user moves about the space, and (3) changes based upon what the user is actually doing. When multiple people are in the same space, each individual’s information will clutter the visual space of other occupants. Unfortunately, the human visual system will automatically attend to nearly all the resultant motion transients, disrupting the feeling of calm.

5 Exploiting change blindness

One way to keep information current without attracting a user’s focus of attention is to exploit *change blindness*. Change blindness is the inability to detect

what should be obvious changes between two images or image sequences. Assume that some large (and often improbable or impossible) change occurs between two otherwise identical images displayed in succession (e.g. a prominent object disappears or a person changes clothing). If a *visual distracter* is displayed either between the image change or precisely at the time of change, people find the thing that changed between the images difficult to detect – even when they are told to watch closely for differences [30, 29]. Once people are alerted to the changes they could not see, they are quite surprised that they could not see the change before.³

Change blindness occurs because a complete mental representation of a scene is never constructed in memory [24]. Even after prolonged viewing, the mental representation of a scene often does not include “obvious” details such as what objects are there, what their properties are, and where they are positioned. Although the brain can instantaneously detect a change (i.e. a motion transient) between two visual stimuli flashed one after another, if the change is obscured either by removing the motion transient or by creating other motion transients, the brain must rely on its sparse mental model of the scene to remember what has changed.

Change blindness can be exploited by the context-aware user interface designer to minimize detectable change within a ubiquitous environment even as digital information is being continuously updated. Users can be made “blind” to changes if the ubiquitous computing systems exploit one of a set of well understood change blindness strategies. The goal is to make changes so that they do not cause detectable motion [14]. Descriptions of some of the strategies follow. In two cases, the strategies have been used in a laboratory ubiquitous environment.

5.1 Blanking an image

Figure 2 shows an example of two images that have been used in change blindness visual memory experiments [22]. Note that Figure 2b has been manipulated so that the traffic line is dashed. When these images are flashed in succession immediately after one another, a viewer will instantly detect the change to the road lines. However, if a blank image is flashed in between the two images for a fraction of a second, experimental subjects typically require anywhere from several seconds to more than a minute of study, where the image is changed about once every second, to detect the line change. It is important to note that *the viewers have difficulty despite being told that a change has definitely occurred* [25].

5.2 Changing views

Film buffs know that many impossible changes between sudden changes in view (i.e. cuts) go unnoticed [11]. Controlled change blindness studies have shown

³ The best way to understand the (sometimes exasperating!) power of the change blindness phenomenon is to view some examples. A set of demonstrations can be found online at <http://www.wjh.harvard.edu/~viscog/change/demolinks.shtml>.



Fig. 2. Change blindness techniques can be used to mask motion transients so that displays can change from (a) to (b) and viewers will have a surprisingly difficult time detecting the change in the road line from solid to dashed [22].

that changes such as switching out actors and clothing type or color are often not detected. In one movie created for a research study, 9 cut-consistency errors (e.g. dinner plates changing color) were made across 5 cuts. On the first viewing, none of ten subjects, *who were told to look for inconsistencies*, detected any. On a second viewing only 2 of the 9 errors on average were noticed [15].

5.3 Displaying “Mud splashes”

Any visual distracter will disrupt peripheral visual processing and permit changes to go undetected. One technique is to “splash” small blocks of pixels on the image for a split second as the two images are switched [22]. Figure 3a shows splashes that have been used with the images in Figure 2. The splashes draw the viewer’s attention, disrupting the visual field. These “mud splashes” do not necessarily cover the region of the image that experiences change, as shown in Figure 3a. Even though viewers instantaneously see the lines disappear when the mud splashes are not used, with the mud splashes naive subjects have a difficult time detecting that the road lines are different.

A ubiquitous environment can exploit this effect by coordinating updates of information displayed in the ubiquitous environment. If a large part of one display will be changed, other displays in the environment can be changed at the same time. Only one motion transient may be perceived by the user.

5.4 Changing information slowly

A change may go unnoticed if the resulting motion transient is masked by the transient from another change (as in the previously discussed techniques). Another way to “mask” a motion transient is to change the image slowly enough so that the change is below the threshold of peripheral detection [31]. Figure 3c shows the first frame of a 12 second video that slowly fades to Figure 3d. Viewers

of this sequence are often unable to detect that any change has taken place in the field in the image even when they study the image intensely looking for a change throughout the sequence. Other studies have shown that entire objects that fade in or out slowly can go undetected as well.




Fig. 3. (a) “Mud splashes” create distracting visual transients and mask the road line change shown in Figure 2 [22] (b) An image of a digital counter displaying the slow-change user interface. (c,d) Slow changes hide motion transients. Subjects intently looking for change typically fail to detect the change from (c) to (d) in a movie that fades over 12 seconds [31].

A version of a slow change information display for an interactive environment has been implemented in our laboratory. Figure 3b shows a digital counter that is displaying a picture that uses slow change to encode personal information for an occupant of the space. The display appears static to the viewer. However, by using slow change to mask motion transients, updated stock data for one company is displayed. The data is embedded in the photograph via computer manipulation of the horizon of the mountain range. Even users who are told that information in the image is changing have a difficult time determining that the horizon is dynamic until the motion is pointed out. Visitors to the environment who are never told there is anything unusual about the image do not notice that dynamic information is being displayed there.

5.5 Exploiting eye blinks or saccades

Eye blinks and saccades also mask motion transients [7, 32]. In future computing environments where occupants are wearing eyepiece displays, blinks could easily be detected. Saccades might also be detectable with eye tracking hardware. With sufficiently fast processing and decision making, changes could be made to digital information during the 100-200 ms blinks or the 20-70 ms saccades without any detectable motion transients. Studies have shown that even changes to text being read can easily go undetected using saccades [19].

5.6 Using occlusion

Occlusion has also been shown to be an effective method for masking motion transients. Occluding objects create motion transients at the leading and following edges as they move. However, these transients are expected. If something formerly in the user's field of view changes while it is in an occluded region, there is no unexpected motion transient to draw the viewer's attention. Substantial changes can go undetected [9].

We have implemented a demonstration using the occlusion change blindness principle. Figure 4a shows the surface of a digital table in our laboratory that can detect a small set of objects placed on it using object silhouettes captured by a camera mounted under the table. In this case, an object is detected in real time. A projector mounted under the table marks the position of the object. Figure 4b shows how the context provided by the tracker and a real-time people tracking system can be used to project a message targeted to one individual at the table. The message will go undetected by the person on the opposite side of the table. When the object is lifted from the table the message disappears and when the object is moved the message is automatically positioned to maintain the occlusion relationship. Not only does one user never see some information that is changing, private information can be conveyed to one without disrupting the second. In addition, private information can be conveyed to a particular user in a public space. If data were being displayed on the digital table, the same system could be used to update that information without creating unnecessary motion transients using the positions of objects on the table as they are naturally moved about.

Figure 5 illustrates how occlusion regions created by moving people could also be used to update information on the walls without creating motion transients. As person-1 moves between person-2 and the data on the wall, the occlusion motion transients can be used to mask changes in data projected on the wall without attracting person-2's attention. When person-2 eventually looks at this information, it will be current. With a sufficient number of people moving about a space, it may be possible to update some slowly-changing information without ever creating an unmasked motion transient. As more people enter a space, however, the geometric reasoning required to do this increases in complexity.

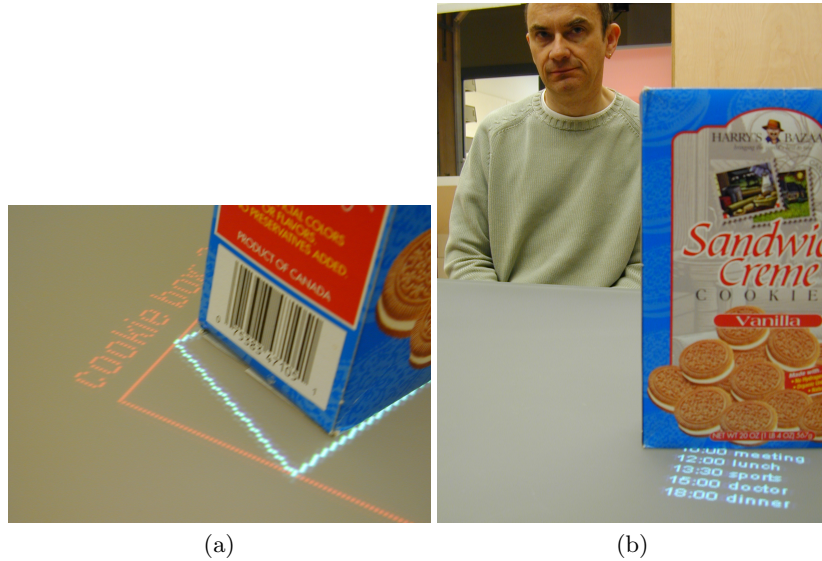


Fig. 4. (a) A digital table in the laboratory that can detect objects placed on top of it and project messages on the table (from the underside of the table). (b) Combined with a digital tracker that knows the location of people, the object tracker can present a personalized message to only one user exploiting occlusion in the environment. As the object is moved around the table, the message stays occluded to those on the other side of the table. The message does not distract the non-recipient. This method could also be used to deliver private information in a public space.

5.7 Other techniques

Several other change-masking methods might be employed by future ubiquitous environments. For instance, laterally shifting an entire image is known to prompt change blindness [4], as is inverting an image. Events that create an intense “startle response” such as an extremely loud noise (or perhaps a phone ringing) may mask otherwise obvious visual motion transients. Finally, change blindness researchers believe that motion transient masking should occur for both visual and auditory signals and have called for studies on “change deafness” [28].

6 A new model for display of information

Although isolated examples can be constructed in the laboratory, ultimately using change blindness in ubiquitous interface design requires context-sensitive applications in the environment to *cooperate* so that users are not confronted with unnecessary motion transients as individual applications update their data displayed in the environment. A shared middleware control mechanism might allow independent interface designers developing for ubiquitous environments




Fig. 5. Context-aware environments that can estimate the position of displays and objects could use occlusion information to change or update information in ways that prevent unexpected motion transients and therefore preserve a feeling of calm.

to present information in a minimally disruptive way.⁴ An appropriate control mechanism might prevent a decline into an information-overloaded environment, even when the parties displaying data have conflicting interests.

Figure 6 shows one possible system. In this scenario, applications no longer directly control precisely when information is updated. Applications instead send *display requests* to a change reasoning control system. Each display request consists of the intended recipients, the priority, the required temporal accuracy, and the display device. These requests are categorized and prioritized. An *activity recognizer* uses sensor data acquired from the environment and other contextual information (e.g. time, knowledge about patterns of behavior) and a user model to select an optimal strategy for updating the display requests in the queue. The *display controller* then sends the updated message to the proper device at precisely the time when motion transients will be minimized. Each application will need to specify the contextual situations in which the information is to be displayed.

For example, if multiple people are in an environment and two applications are occasionally updating the data of two different users, the change reasoning system would accept display requests from each application. Instead of simply using two ambient displays that update independently, both displays whenever possible will update at the same instant, thereby minimizing detectable motion transient events. Further, whenever possible these changes will be correlated with other disruptors since only a single “splash” will typically be detected.

Such a system requires individual user interface designers to relinquish some control over exactly when data displays are updated. How appropriate this type of model depends upon the application. A video game that requires precision updating would only submit change requests that ignored context and used immediate updating. However, applications that do not require precision timing (e.g. a puzzle game, certain types of reminders, educational applications, shared

⁴ Context-aware sensing that uses sensor fusion from sensors distributed in different products is also likely to require middleware software (e.g. [27]).

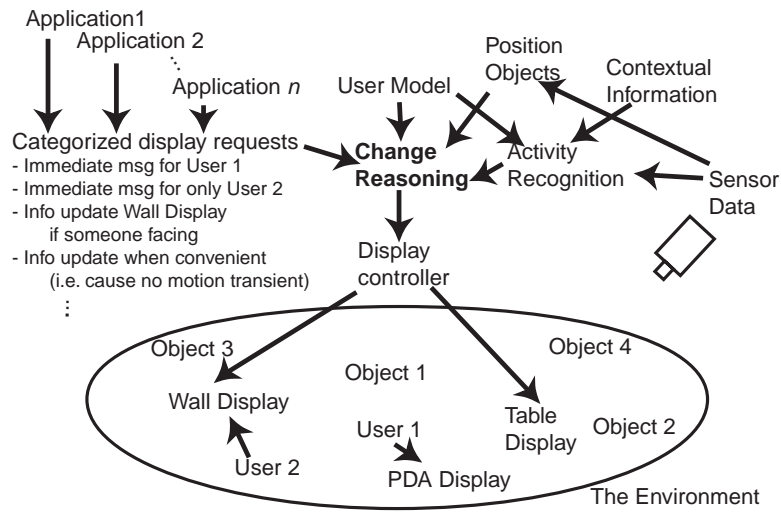


Fig. 6. This diagram shows how a change reasoning system might mediate the display in a ubiquitous computing environment so as to exploit change blindness strategies in order to minimize detectable motion transients and maximize a feeling of calm.

awareness of remote spaces) are appropriate for the model. Designers would lose precise control over when and how messages were conveyed but gain the ability to write software that works in tandem with other ubiquitous computing interfaces and maintains a calm environment. Is the tradeoff worth it? The detailed operation of such a system and construction of a prototype is left as future work once context-aware environments with ubiquitous displays are more widely available for robust user interface studies.

7 Limitations

Using change blind user interface design strategies has limitations. Inevitably some messages will create motion transients, and to fully exploit the technique will require robust object and people tracking.

Most challenging is that some masking methods are not effective when the changes occur to the object of *central interest* in the scene (e.g. a face of a key person in an image) as opposed to an object of *marginal interest*. Change detection time for central interest objects is fast regardless of object color, position, and presence/absence [25]. Object position and presence are better encoded by the brain than surface properties, which makes these properties more difficult to change without triggering a detectable motion transient [2]. Therefore, change blindness is less likely to be effective for objects of strong interest. The worst-case scenario, however, is no worse than the current situation: a motion transient is created that mildly attracts the user's attention. Detecting what may be a user's

central interest versus marginal interest in a ubiquitous computing environment is an active research question.

Fortunately, some clues as to the likelihood that a masked change will be detected are available. For instance, user identification of a masked object removal is more likely than an object addition. Identification of a masked addition has been shown to be as likely as identification of a color change. Location changes are significantly less likely to be identified than addition and color changes [20]. Masked deletion of a unique object is easier to identify than addition of an object not present in the first scene [3]. Cognitive scientists have not yet tested change blindness in truly dynamic scenes and explored detection of changes to object velocity and acceleration fully [34]. Finally, recognition and recall improve with longer viewing [33]. All of these properties could be taken into account by the change blindness reasoning module.

It is important to note that in change blindness studies subjects are typically told that changes will occur and yet often have great difficulty seeing them. In a ubiquitous computing environment, subjects will not be focussed on detecting changes.

8 A solution in search of a problem?

Given that we do not currently have home environments with rich display capability, is this a solution in search of a problem? For instance, if there are too many ubiquitous surfaces each displaying too much information, then will it not be possible to either reduce the number of surfaces or to have them display less?

The practical reality is no. The user interfaces of the future will be a collection of devices created by a collection of entities. Context-aware sensing and user interface design will be valuable because information can be proactively presented to people when and where they want and need it. If the future interfaces create such a digitally cluttered environment that home occupants must turn off displays to get a bit of calm, the enormous potential of ubiquitous computing will never be realized.

Now is the time for researchers to consider the serious problem of how to present context-aware information ubiquitously in home environments while preserving the aesthetic properties of the home that make it a desirable place to be. A system like the one proposed in Section 6 that exploits the powerful phenomena of change blindness may be more effective at preserving this aesthetic than a proliferation of uncoordinated ambient displays.

9 Summary

Change blind user interface design has the potential to allow for the creation of future ubiquitous computing environments that minimize detectable motion transients and preserve the desirable feel of a calm environment. Exploiting the technique may require user interface designers to relinquish some control over

exactly how and when information is displayed. In return, occupants of the environments may be able to enjoy the benefits of context-sensitive messages without the uncomfortable feeling that they are being inundated with distracting, changing digital information.

Acknowledgments

Joachim Böttger implemented the digital table tracking system and the slow-change mountain data display. Chuck Kukla helped to clarify the presentation of ideas. Byron Stigge built the digital table. Leonardo Bonanni built the digital counter. Figure 1a was rendered by Rujura Hongladaromp.

References

1. G.D. Abowd and E.D. Mynatt. Charting past, present, and future research in ubiquitous computing. *ACM Transactions on Computer-Human Interaction*, 7(1):29–58, 2000.
2. V. Aginsky and M.J. Tarr. How are different properties of a scene encoded in visual memory? *Visual Cognition*, 7(1/2/3):147–162, 2000.
3. G. Agostinelli, S.J. Sherman, R.H. Fazio, and E.S. Hearst. Detecting and identifying change: Additions versus deletions. *Journal of Experimental Psychology: Human Perception and Performance*, 12(4):445–454, 1986.
4. S.J. Blackmore, G. Brelstaff, K. Nelson, and T. Troscianko. Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception*, 24:1075–1081, 1995.
5. J.S. Brown and P. Duguid. Keeping it simple. In Terry Winograd, editor, *Bringing Design to Software*. Addison-Wesley Pub Co, 1995.
6. B. Buxton. Integrating the periphery and context: A new taxonomy of telematics. In *Proceedings of Graphics Interface*, pages 239–246, 1995.
7. J. Grimes. On the failure to detect changes in scenes across saccades. In K. Akins, editor, *Vancouver Studies in Cognitive Science*, volume 2: Perception. Oxford University Press, 1996.
8. J.M. Heiner, S.E. Hudson, and K. Tanaka. The Information Percolator: Ambient information in a decorative object. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 141–148. ACM Press, November 1999.
9. J.M. Henderson. Transsaccadic memory and integration during real-world object perception. *Psychonomic Science*, 8:51–55, 1997.
10. E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse. The Lumiere Project: Bayesian user modeling for inferring the goals and needs of software users. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, pages 256–265, San Francisco, July 1998. Morgan Kaufmann.
11. J. Hotchberg. Representation of motion and space in video and cinematic displays. In K.R. Boff, L.K. Kaufman, and J.P. Thomas, editors, *Handbook of perception and human performance*, pages 22.1–22.64, London, 1986. John Wiley & Sons.
12. G. Humphreys and P. Hanrahan. A distributed graphics system for large tiled displays. In *Proceedings of IEEE Visualization '99*, October 1999.
13. N. Jeremijenko. The dangling string. Artistic Exhibit.

14. R. Klein, A. Kingstone, and A. Pontefract. Orienting of visual attention. In K. Rayner, editor, *Eye movements and visual cognition: Scene perception and reading*, pages 46–65, New York, 1992. Springer.
15. D.T. Levin and D.J. Simons. Perceiving stability in a changing world: Combining shots and integrating views in motion pictures and the real world. *Media Psychology*, 2:357–380, 2000.
16. B. MacIntyre, E.D. Mynatt, S. Voida, K.M. Hansen, J. Tullio, and G.M. Corso. Support for multitasking and background awareness using interactive peripheral displays. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology*, pages 41–50, New York, NY, 2001. ACM Press.
17. P.P. Maglio and C.S. Campbell. Tradeoffs in displaying peripheral information. In *Proceedings of the CHI 2000 Conference on Human Factors in Computing Systems*, pages 241–248, New York, NY, 2000. ACM Press.
18. C.C. Mann. Electronic paper turns the page. *Technology Review*, March 2001.
19. G. McConkie and D. Zola. Is visual information integrated across successive fixations in reading? *Perception and Psychophysics*, 25:221–224, 1979.
20. S. Mondy and V. Coltheart. Detection and identification of change in naturalistic scenes. *Visual Cognition*, 7(1/2/3):281–296, 2000.
21. D.A. Norman. *The Design of Everyday Things*. Doubleday, New York, NY, 1989.
22. J.K. O’Regan, R.A. Rensink, and J.J. Clark. Change-blindness as a result of “mudsplashes”. *Nature*, 398:34, 1999.
23. C. Pinhanez. The Everywhere Displays Projector: A device to create ubiquitous graphical interfaces. In G.D. Abowd, B. Brumitt, and S.A.N. Shafer, editors, *Proceedings of the Conference on Ubiquitous Computing*, LNCS 2201, pages 315–331, Berlin Heidelberg, September 2001. Springer-Verlag.
24. R.A. Rensink. The dynamic representation of scenes. *Visual Cognition*, 7(1/2/3):17–42, 2000.
25. R.A. Rensink, J.K. O’Regan, and J.J. Clark. To see or not to see: the need for attention to perceived changes in scenes. *Psychological Science*, 8:368–373, 1997.
26. J.W. Rohrbaugh. The orienting reflex in humans. In *Varieties of Attention*, pages 323–373. Orlando Academic Press, 1984.
27. D. Salber, A.K. Dey, and G.D. Abowd. The Context Toolkit: Aiding the development of context-enabled applications. In *Proceeding of the CHI 99 Conference on Human Factors in Computing Systems : the CHI is the Limit*, pages 434–441, New York, NY, 1999. ACM Press.
28. K.L. Shapiro. Change blindness: theory or paradigm? *Visual Cognition*, 7(1/2/3):83–92, 2000.
29. D.J. Simons. Current approaches to change blindness. *Visual Cognition*, 7(1/2/3):1–15, 2000.
30. D.J. Simons and D.T. Levin. Change blindness. *Trends in Cognitive Science*, 1(7):261–267, 1997.
31. D.J. Simons, R.L. Reimer, and S.L. Franconeri. The role of abstract representations and motion signals in change detection. In *Proceedings of the Association for Research in Vision and Ophthalmology*, April 2000.
32. W.S. Sun, R.S. Baker, J.S. Chuke, B.R. Rouholiman, S.A. Hansan, W. Gaza, M.W. Stava, and J.D. Porter. Age related changes in human blinks. *Investigative Ophthalmology and Visual Science*, 38(1):92–99, 1997.
33. B. Tversky and T. Sherman. Picture memory improves with longer on time and off time. *Journal of Experimental Psychology: Human Learning and Memory*, 104:114–118, 1975.

34. G. Wallis and H. Bühlhoff. What's scene and not seen: Influences of movement and task upon what we see. *Visual Cognition*, 7(1/2/3):175–190, 2000.
35. M. Weiser and J. Brown. Designing calm technology. *PowerGrid Journal*, 1(1), 1996.
36. Craig Wisneski, Hiroshi Ishii, Andrew Dahley, Matt Gorbet, Scott Brave, Brygg Ullmer, and Paul Yarin. Ambient displays: Turning architectural space into an interface between people and digital information. In N.A. Streitz, S. Konomi, and H-J. Burkhardt, editors, *Proceedings of Cooperative Buildings, Integrating Information, Organization and Architecture, First International Workshop, CoBuild'98*, LNCS 1370, pages 22–32. Springer, 1998.