

Desituating Action: Digital Representation of Context

Jonathan Grudin

Microsoft Research

RUNNING HEAD: DESITUATING ACTION

Jonathan Grudin, Microsoft Research, One Microsoft Way, Redmond, WA,
98052 jgrudin@microsoft.com

Jonathan Grudin is a Senior Researcher in the Adaptive Systems and Interaction Group at Microsoft Research, focusing on computer supported cooperative work and the adoption and use of technologies supporting collaboration.

ABSTRACT

Many psychological studies have shown that when we act, and especially when we interact, we consciously and unconsciously attend to context of many types. Sensors can pick up some but not all context that is acquired through our senses. Some context is lost, some is added, and captured context is presented in new ways. Digital aggregators and interpreters do not aggregate and interpret the same way we do. Missing or altered context disrupts our processing of information in ways that we may not recognize. To address the disruption we may use additional sensors to capture and deliver some of the missing context. Learning to handle these new conduits is then a further source of disruption, and on it can go. With greater knowledge of context, we can work and interact more efficiently, assuming that we can take learn to take advantage of the information without being overwhelmed. However, converting contextual information to a digital format changes it in specific ways. Transient information becomes more permanent, local information is made available globally, and information that once spread slowly can spread much more quickly. The information can enable us to work more efficiently, but these changes in its nature have profound indirect effects. The potential loss of privacy is widely discussed, but other effects may be more significant. In particular, the loss of confinement and transience of information creates an environment that is fundamentally unnatural, in conflict with the one we evolved to live in.

CONTENTS

1. INTRODUCTION
2. DISCUSSION OF THE DEY ET AL. FRAMEWORK
 - 2.1. A Major Contribution
 - 2.2. Benefit and Cost of Simple Applications
 - 2.3. Defining Context: Theory And Practice
 - 2.4. Context and Communication
 - 2.5. Context When Location-Awareness Is Perfect
 - 2.6. Two Scaling Issues
 - 2.7. When Is Some Context Better Than None?
3. ENHANCING EFFICIENCY
4. THE TRANSFORMATION OF CONTEXT
5. FENCING IN THE LAST FRONTIER
 - 5.1. Privacy: Concerns May Be Misdirected
 - 5.2. The Separation Of Action From Time And Space
 - 5.3. Desituated Action: When Local Becomes Global
 - 5.4. Informal Interaction: Endangered or Irrepressible?
6. CONCLUSION

1. INTRODUCTION

In this essay I describe a key challenge for the context-aware applications outlined by Dey, Salber and Abowd (2001 [this special issue]): Recognizing and communicating context is a path to more efficient and effective use of technology, but capturing context digitally alters it fundamentally. The context that is captured is removed from *its* context, namely the context that is not captured. Furthermore, aggregation or interpretation that is done by software differs from aggregation and interpretation that are done by biological, psychological, and social processes. Finally, once digital information reaches a network it could appear anywhere on the planet at any future time. Consequently we find ourselves

in a new situation, a new context. Until we adjust, it is almost inevitably disruptive. Sensible reactions to the disruption range from capturing yet more context to creating security features that block or filter contextual information. The former continues the cycle of disruption and compensation, the latter can add substantial overhead and undermine the benefits accrued by capturing context in the first place.

In Section 2 I frame the discussion. Dey et al. define context broadly but then focus on a narrow range of applications. They postpone consideration of broader aspects of context, scaling and indirect effects. Sections 3 and 4 introduce the key point of this essay, that representation of context can boost efficiency, but fundamental changes occur when information is made digital. Section 5 explores and illustrates the consequences of this transformation.

2. DISCUSSION OF THE DEY ET AL. FRAMEWORK

2.1. A Major Contribution

The authors have done a great service by setting out a lucid framework for context-aware applications and by developing that framework through the design and construction of five simple applications. The paper makes a compelling logical argument for the elements of their architecture. The enterprise of context-aware applications becomes very intellectually appealing. The analogy of the evolution of graphical user interfaces seems helpful. In addition, they describe their tools being used by people not involved in the tool development, a nice illustration of human-centered tool design.

2.2. Benefit And Cost Of Simple Applications

The applications that helped develop the ideas and tools were used in a university research setting. This seems the right approach for this very complex research topic. Iterative design of applications and tools is essential to discovering what is needed. A full-fledged real-world application would take too much time to allow sufficient iteration. By gradually tackling more complex applications the authors can make methodical progress. They realize that in scaling up to address real-world problems, their architecture will evolve. Whether this occurs through the addition of layers or through more fundamental redesign, the kind of experimentation described here will identify requirements for future context-aware application architectures.

The price paid for relying on simple prototypes in research settings is well known: This is not user-centered design. Dey et al. argue convincingly that, similar to the inventors of GUIs, they have identified and addressed a crucial issue. However, despite exemplary research lab testing, the first GUI inventors were not commercially successful. They attributed this to factors that include a lack of understanding of their customers and an architecture that did not scale in various ways (Johnson et al., 1989). Until real

applications are tried, we do not know what is involved in making them work, or even whether or not they are workable.¹

2.3. Defining Context: Theory And Practice

The paper does a nice job of distilling a definition of context from the literature, concluding that “context is typically the location, identity, and state of people, groups, and objects” where state refers to “physical, social, emotional or informational state.” In practice, however, the authors restrict themselves to location and identity. This bias is pervasive. For example, the active badge call-forwarding application of Want et al. (1992) is described as “arguably the first context-aware application.” This ignores, for example, prior work on the media space concept of “What You See is What I See” (WYSIWYS), which posited a need to preserve context in communication systems. The authors are aware of their restricted focus, concluding that there is a need for “moving beyond the environmental context used in traditional context-aware research.” However, their five applications focus on location and identity.

I had to choose whether to consider context in terms of the broad definition or to restrict my remarks to location, presence, and sensor-based context. The essay focuses on location-aware systems, but the challenges are even greater when we go beyond environmental context to include “emotional and information states.”

2.4. Context And Communication

Another important distinction contrasts what are essentially individual productivity tools with systems in which interpersonal interaction is central. Context issues scale up dramatically when communication among people is involved (as illustrated by the WYSIWIS work). Location-based work, including that of Dey et al., focuses primarily on individual productivity. Even when locating people is the point, the interaction is handled outside the context-aware application.

A location-aware museum guide is a purely individual productivity system. A system that allows one to link notes to a recorded lecture involves multiple people, but the primary use is for individual productivity, not communication. Similarly, an automated In/Out board is primarily an efficient tool for individuals, it does not directly support much communication. It supports one ‘bit’ of communication, not insignificant, but dramatically different from, say, a videoconferencing application. Videoconferencing

¹ Consider the most complex prototype, the Conference Assistant. It is a valuable design exercise, but is it a viable application? Will people enter keywords hoping that a system will identify interesting sessions, or prefer to skim the program or proceedings? Keyword matching is unreliable, and people often use criteria other than content: Some prefer panels, others avoid panels; many are guided by a speaker’s notoriety or track record, and many make their decision based on what else is available at the time. Notes can be linked to recorded audio, video, and slides, but few conferences record any of these, and for a range of social factors this seems unlikely to change. Would you make the effort to continually indicate your current level of engagement on the chance that a colleague is looking for a session to join? How comfortable would you be knowing that there is a public record of when you slipped in and out of a presentation? Such questions suggest that a field study of conference attendance would be valuable prior to full implementation of this system.

introduces a raft of additional context issues: What are the other people talking about, what are they doing, what do they already know about the matter at hand, what documents do they have access to, what is their emotional state, how are they reacting to what is being said, where is their gaze really directed, and so forth.

By focusing on individual productivity tools, Dey et al. took a reasonable research approach: Context problems arising from interaction are the focus of other researchers. In considering this work, keep in mind that issues surrounding interaction must also be addressed if we consider context in its broadest sense.

2.5. Context When Location-Awareness Is Perfect

The role of environmental or location awareness in the broad picture of context is revealed when we have perfect or near-perfect awareness of physical context, as in non-mobile environments such as desktop computer use. Location is known, other context problems remain, affecting both individual and group use of technology.

Temporal context is lost: Applications do not generally record or make much use of a person's history. Attentional context is missed: Applications do not know where a user's attention is focused, whether he or she is busy or open to interruption, talking to someone or on the phone, and so forth. User modeling has not been very successful at representing "physical, social, emotional or informational states."

Progress has also been slow in supporting small groups of desktop system users. At the first Computer Supported Cooperative Work conference, Trigg, Suchman, and Halasz (1986) noted the lack of software support for "mutual intelligibility," the communication of context. In an echo appearing over a decade later, Erickson and Kellogg (2000) claimed that software applications are "socially blind," despite the efforts of many researchers and developers in the intervening years to address the problem.

Organizational information systems are also lacking. A visitor or new employee rarely finds tools that provide useful context, such as good online building maps, guides to organizational structure and document repositories, and in-house expertise locators. Dey et al. call for capturing impromptu meetings, but rarely can records be accessed after an event even for scheduled meetings in reserved conference rooms.

In conclusion, many context issues remain when location is known.

2.6. Two Scaling Issues

The authors note that wide-scale deployment of context-aware applications will require vast arrays of sensors, widgets, interpreters, aggregators, discoverers, and services. It will involve extensive relaying of information and require the recording of comprehensive histories of activity of all components and interactions. Scaling up will raise technical issues and require higher-level abstractions.

Rapid progress is likely in addressing the technical issues. But how useful will the structure emerging from this bottom-up effort be? Consider the analogy of an ant colony, with its many basic sensors that detect water, food, temperature, other ants, threats, and so forth. The colony exhibits aggregation, interpretation, and discovery. It works beautifully. But its repertoire is quite limited, and it took millions of years of trial and error to arrive at an architecture to accomplish that much.

Top-down approaches to context problems are illustrated by the Cyc project (Lenat, 1995). Cyc, intended to be a repository of “common-sense” information and reasoning that spans all domains of knowledge, is another architecture that promotes the reuse of contextual information by diverse applications. Although Cyc has been criticized on grounds ranging from technical to ontological, its basic rationale is compelling and similar to that of the context-aware framework. Perhaps bottom-up architectures for context-aware applications that seek higher-level abstractions will merge with top-down efforts such as Cyc. But it is not clear that Cyc will deliver on its promise (originally forecast to take ten years, Cyc is almost 20 and its anticipated size has grown an order of magnitude). On the whole, partial implementations of Cyc are not useful. Will the same be true for the equally ambitious vision of context-aware applications?

2.7. When Is Some Context Better Than None?

“You are only about six acquaintances away from anyone else in the world.” “A butterfly flaps its wings in Beijing and a month later a typhoon strikes Bali.” Such familiar claims point to the extreme interrelatedness of people and events, the density of the contexts in which we function.

Psychological studies have demonstrated that people respond *unconsciously* to shifts in context at all levels—sensation, perception, cognition, emotion, social. Human beings make handling context look easy. This can lead us to underestimate the subtleties and complexities of digitally acquiring, representing, and acting on contextual information.

Dey et al. note that considerable information is required to establish useful context. They write, “a component that acquires context information should maintain a history of all the context it obtains.” Even if there is no pending request for the information, “in the future... an application may need (it)... the architecture must support the storage of context.” But we can’t represent, much less store, it all. Anyone’s actions could affect us, but we can’t yet represent 6 billion people. Anticipating tomorrow’s weather possibilities could be useful, but we can’t track all Chinese butterflies.

In what circumstances can we collect enough context to be useful? When will the value of the available information outweigh problems caused by the distortion introduced by partial representation of context? This is explored in the sections that follow.

3. ENHANCING EFFICIENCY

Digital technologies will, I believe, enhance safety, improve accountability, and provide more equal access to resources and opportunities. There is a growing consensus

that efficiency gains from digital technology contribute to rising productivity. Providing additional context digitally can increase efficiency by providing context in online settings that is normally available in physical settings.

When we interact face to face, I can tell at a glance who else is present, how you are reacting to what is said, what objects are around, and so forth. When we interact online, explaining all this verbally is inefficient, but if it is not explained, confusion often results. The solution is to convey more of the context automatically.

The digital representation of contextual information can also help by obviating the need to track down the information in other ways. A location-aware museum guide is more efficient than a pamphlet one must search through. A route finder that senses the location of your vehicle saves you from working out a path; if you can also avoid a congested highway it is even more efficient. An automatically recorded meeting eliminates the need to take detailed notes or consult with someone who attended. If it provides a more accurate record of what was done than would otherwise be available, this can increase efficiency as well.

Even when a new technology will increase efficiency, it often disrupts existing practice initially. Disruptions can be addressed in three ways: (i) the technology can be rejected; (ii) over time new practices and conventions governing its use can emerge that reduce the inefficiencies; (iii) the technology can be enhanced to compensate for the problems, often by delivering additional contextual information. Even when the benefits of a new technology already compensate for the disruption, (ii) and (iii) can amplify the benefit.

The following examples illustrate the patterns described in (ii) and (iii):

- A real-time application sharing technology enables a group to meet virtually, saving travel time. However it results in confusion at first: Participants are uncertain who else is present, what they are referring to, and so forth (Mark et al., 1999). This is handled by developing social conventions such as announcing arrivals and departures, or technologies that increase the shared context such as videoconferencing.
- Intercom, the Dey et al. call forwarding system, automatically routes calls to the phone nearest to you, enabling callers to reach you more efficiently. However, it can create misunderstanding or confusion when a caller assumes that you are in your office. Providing additional context by letting callers know where you are and who is with you eliminates this source of confusion.
- Presenters who use a desktop video telepresentation system miss seeing audience reactions, so designers build in features such as chat channels and emotion or applause meters, which presenters then must learn to handle (Begeman et al., 1986; Isaacs et al., 1994).

The desire for greater efficiency presses us toward ever greater representation of context. The context-aware application architecture of Dey et al. is motivated by this inexorable drive. Consider again Intercom, which automatically routes calls to the phone nearest Sam unless someone else is in the room, in which case it first discretely inquires. Efficient, but potentially embarrassing if someone is standing in the hall by an open door, within earshot. To deal with this, we could extend the system to include sensors monitoring the doors to the room (are they open or closed?) and detecting people within earshot outside open doors. If we model individual differences in auditory acuity we could do even better. In such ways, representing some context digitally can lead to pressure to provide and use yet more digital context.

4. THE TRANSFORMATION OF CONTEXT

Putting contextual information on a digital network fundamentally transforms it. It can lead to relatively subtle distortions as well as to very substantial shifts in use.

The physical world presents a dense multi-dimensional contextual space. Different aspects of context assume different levels of importance, and these can vary over time. Technology designers must determine which contextual threads to capture and where to place the significance cutoff points. These decisions are based partly on the importance of a particular type of information, and partly on what can be captured or inferred easily or economically. Context that is expensive to capture must have a big payoff, context that is relatively easy to capture only requires a small payoff.

When both salience and ease of capture determine what context is represented, there can be an unevenness across context threads. In one case only highly salient information is provided (because capture is expensive), which could cause people to overrate the salience of other information that is available mainly because providing it is easy. Conversely, when a lot of relatively unimportant context is provided because it is inexpensive, people may fail to realize that other, expensive context is missing.

Anything that does not “make the cut,” (i.e., is not captured) is invisible to someone inspecting the digital record at a different location or time. Anything that is recorded instantly achieves a potential pervasiveness and immortality that it did not have before. A word spoken in an unrecorded meeting is heard only in that room and is gone once the echo subsides. When sucked digitally into a network it can potentially be heard anywhere in the world, in real time or at any future time. This profoundly changes the nature of that information and the contexts created from it.

A well-known problem with the quasi-immortality of data that reaches the network is that erroneous information can be impossible to fully correct. A more subtle problem is that contextual information that is represented digitally is never completely accurate, because some of the related context is not included. For example, when New Jersey moved to publicize and seek comment on planned prisoner paroles, opponents argued that members of the public “would not necessarily have all the background that the parole board does... a whole file of information about education programs, work release and other programs the inmate had been involved in, information not included in a six-line

description on the Web site.” (New York Times, 7 November 1996). This could arguably be addressed by providing more context online (the inexorable tendency noted above), but would people read it? Another example: a conference program committee made all reviews of all papers visible to the full committee via the Internet to make deliberations more efficient. One person complained that the review was now taken out of context: Instead of being read only by people familiar with the paper, it would be seen by people who had not read the paper.

Learning to work in a world of increasingly context-aware applications is one of the greatest challenges that we face. Efficiency gains *can* be realized by representing and communicating information digitally. Increasing the available contextual information can help, but raises concerns about privacy, accuracy, and interpretation. The genie is out of the bottle, loose in a world increasingly committed to efficiency. We must learn to live comfortably with it. The next section explores and illustrates key challenges.

5. FENCING IN THE LAST FRONTIER

5.1. Privacy: Concerns May Be Misdirected

Direct marketers argue that they can more efficiently match supply and demand by collecting personal information—our neighborhood, purchases, hobbies, medical records, family additions, and so forth. Efficiency can conflict with a desire for privacy. Loss of privacy is often seen as the primary risk in embracing ubiquitous computing.

Feelings about privacy are important. Watchdog groups, legislators and executives, European governments, standards organizations, and many researchers are focusing on privacy issues. Dey et al. suggest that privacy concerns can be addressed by assigning to each widget an owner who establishes rules governing access. This is not a serious suggestion, given the vast numbers of widgets and the fact that application-independence will enable each to be queried for purposes unimagined by the owner. Everyone would become a full-time widget maintainer. For example, the decidedly not deaf DUMBBO mobile meeting board prototype, which automatically detects and records conversations carried out in its vicinity, has a location widget “for each possible location of the mobile board.” In the unlikely eventuality that people would agree to manage the many location-specific widgets, privacy concerns could still foster a resistance to DUMBBO.

Another example: To foster informal interaction, audio and video was used to link three public areas in two buildings occupied by one corporate division (Jancke et al., 2001). Although these were public spaces and viewing was largely reciprocal, 5% to 10% of employees felt the system violated their privacy and two or three objected violently. Privacy concerns appear to be growing.

Nevertheless, privacy may not be a major issue in the long run. Whether or not Scott McNealy’s recent widely-publicized dictum “You already have zero privacy. Get over

it,”² reflects the current reality, we are moving in that direction, as outlined in the provocative book *The Transparent Society* (Brin, 1998). Privacy is a relatively new concept and, as currently conceived in developed countries, a relatively new phenomenon. Through most of our existence, human beings lived in small groups or single-room family dwellings, and worked in open spaces.

When people see benefits that outweigh risks, they voluntarily adjust their comfort levels regarding privacy. Surveillance cameras are accepted in banks, stores, and increasingly in outdoor public places to reduce crime. (In the Jancke et al. study, employees upset at the cameras deployed to foster informal interaction shrugged off hidden surveillance cameras deployed in the building for security.) Despite well-circulated accounts of the extensive collection, aggregation, and interpretation of credit and debit card transactions to identify purchasing patterns, people would rather use them than make the effort to carry cash. In some organizations, employees cheerfully allow open access to details of their online calendars because it is extremely efficient for scheduling meetings and a range of other purposes (Palen, 1999). Teams that are initially skeptical about working in large open spaces (“war rooms”) change their views after experiencing improved efficiency in communication and collaboration (Teasley et al., 2000). In my company, individual offices are the rule, but telephone support centers adopt open office plans in order to more easily share expertise in response to calls.

In general, people adjust by redefining privacy and by establishing new practices and social protocols (Friedman, 1997; Murphy, 1964; Palen, 1999). Just as people who live in buildings with paper-thin walls may adopt a convention of ignoring what they cannot help overhearing, people who allow open access to their calendar details assume that people will access information only when needed and would be offended by an inquiry that revealed “snooping.” Being able to block off a calendar entry or reserve a conference room is deemed an adequate balance. Privacy is ultimately a psychological construct, with malleable ties to specific objective conditions.

Why then the uneasiness, the widespread attention to privacy? It may reflect an awareness at some level of something more fundamental than privacy that is being challenged: The steady erosion of clearly situated action. We are losing control and knowledge of the consequences of our actions, because if what we do is represented digitally, it can appear anywhere and at any time in the future. We no longer control access to anything we disclose.

In a U.S. job interview, a prospective employee would be offended if an employer asked “Do you have children?” Yet the same employee might at some point ask about schools, and once employed discuss family freely. It is not privacy as much as control of the circumstances of disclosure that is at issue. Of course, if the prospective employee has ever discussed family matters on a newsgroup, the employer could potentially access the information prior to the interview. It is control of disclosure that is going away.

² Illustrating the problem of accuracy of digital information, McNealy’s precise words may be lost to history: Typing his name and “privacy” into a search engine turned up 5 variations in wording, all in quotation marks. Many also appear if his name is misspelled “MacNealy” or “McNeally.”

5.2. The Separation Of Action From Time And Space

If people had less privacy in pre-history than most of us do today, they experienced something we do not: immediacy. Almost every action was a response to a threat or opportunity directly before them. In the words of Masanao Toda (1999), people in the wilderness state acted in the ‘here and now.’ Our emotions and urges—anger, empathy, fear, desire, establishment of social status, and so on—evolved to function well in situations in which people reacted to events in front of them. “Real-time and collocated,” we might say. Advanced agriculture led to larger social units in which social status must be worked out across time and space, in which anger or fear are often triggered by distant events, and so forth. Toda argues that over the past several thousand years we worked out, with only mixed success, societal approaches to social control that channel emotions and urges that are often maladapted to the “asynchronous, distributed” situations we are now often in.³

The distancing from the immediacy of cause and effect has been increased steadily by technologies such as transportation, but above all by communication technologies. Writing, printing, telegraph, telephone, film, radio, television—all allow us to address people across time, space, or both. In politics, for example, oratory and response to the local audience were once central. With television it is riskier to tailor comments to a local audience (comments can be relayed anywhere), and highly edited multimedia productions are replacing speechmaking in national campaigns. Immediacy gives way to a focus not anchored to a single place or moment.

The Internet and Web exemplify the communication and storage technologies that are completing the rupture of action from time and place. When captured digitally and put on a network, a remark or event can show up any place, anytime. It may not, but it *could*. So we never truly know the full context in which we act—where, when, or by whom our actions may be interpreted. Casual email can turn up in a court case years later, for example.

Context-aware platforms make available information from the past and present to help inform our current activity; ironically they also insure that we cannot grasp the full geographical and temporal contexts of our current actions. The next section illustrates and discusses this phenomenon.

5.3. Desituated Action: When Local Becomes Global

Several years ago, researchers at Bellcore built a system to archive and search newsgroups, in order to access contributions by people they admired. To test it they entered the names of colleagues. “We found,” one recounted, “that we were discovering things about our colleagues that we didn’t want to know.” (Jim Hollan, personal communication, 1992.)

³ It is probably no coincidence that the hugely popular horror and action film genres, live and televised sporting events, and video games place viewers in ‘here and now’ situations that we are physically built for but sheltered from in most daily life.

Prior to Deja News, and to some extent still today, newsgroup discussions felt like conversations in a room, ephemeral chatter among like-minded people. But what is written could later be read outside of that context, by anyone, anytime, anywhere. Those who post to a newsgroup or create a web page, anyone with any information from or about themselves on a network, may think that they know when, where, and by whom it will be accessed, but they cannot be sure. The likelihood of information reaching different parties varies, but we can never be confident that we are operating in a local context. Long ago, people *always* operated in a local, situated context. To a large degree we usually still assume that we are.

A major change wrought by digital media is that when past indiscretions were recorded, they were unlikely to surface. This is no longer true. Consider these two cases.

“I went on the job because one of the two models got sick, and the other one was Claudia Schiffer. Obviously I was kind of the nobody, so she got all the work—it was a week in the Seychelles, and they were like, Well, we'll shoot a couple covers on you when we get a chance. I was 17, maybe 18, and they asked me to take off my top for one of these pictures.” — Amber Valleta, model and actress, *New York Times*, June 25, 2000.

“It just brought up a lot of questions, because now that I'm older and married and I have a 14-year-old stepdaughter, and she sees me naked on a Web site, I'm like: How am I going to respond to her?” — Rosie Perez, actress, *New York Times*, June 25, 2000.

Two unknown young women allowed photographs, assuming that their distribution would be limited. When the women became successful, the photos were commercially exploitable. This could have happened prior to the Web, but family members and friends would probably not encounter one published in a salacious magazine, and the magazine issue would soon be difficult to find. Now and forever, anyone with Internet access can stumble across the pictures by typing the actress's name into a search engine, an action that a well-meaning friend, family member or fan might do. As more information is represented digitally, the lives of less famous people will be similarly accessible.⁴

A Ph.D. student approaching graduation and the job market wrote “Thinking about that reminds me of numerous discussions I've been having with grad students and junior faculty about presentation of self on the WWW. There's a real dilemma for these people about how much to 'personalize' their home pages. They're concerned that presenting self means offending others or worse yet projecting bad self image to recruiters or peers so recently (at least among the people I spoke with) a lot of folk have begun de-personalizing their pages. I've thought this through and put a disclaimer at the top of my page as I don't really want to lose the intimacy of putting things on my page that I think are fun... although I am still debating removing a fairly strong political statement...” (Rebecca Grinter, email message, 20 August 1995)

⁴ On 8 January 2001, the online *New York Times* published a list of movies put on an amazon.com site, prior to his notoriety, by an accused mass murderer, along with an analysis of his comments posted on newsgroups. The film ‘Nightmare Before Christmas’ was taken as evidence of an attraction to violence, the journalist apparently unaware that it is a gentle comedy.

This observation made early in the life of the Web was on target. Not long afterwards I was on a hiring committee that rejected a candidate because the committee chair felt “her web page isn’t professional.” She was evaluated on her current web page. Today, services are archiving large regions of the Web. Soon, prospective employers and others may be able to access old web pages—graduate student pages written for friends and other grad students, unpolished undergraduate pages, even embarrassing high school statements.

Brin (1998) describes a lawsuit seeking a computer-based FBI compilation of publicly available information from courthouses across the U.S. The Supreme Court ultimately backed the FBI’s right to withhold it, arguing that although the information was public, citizens had benefited from a “practical obscurity” because only the police or wealthy organizations had the resources to collect it all. Of course, once the dispersed public information is made available digitally anyone could compile it without having to ask the FBI.

In some cases, the “practical obscurity” that parolees and actresses once had only has to last a few hours to be useful. An Edupage summary of a Financial Times article (5 December 1995) titled “Tokyo Exchange Says Internet Is Too Fast” reveals a change in context wrought by the speed of information dissemination: “Because of ‘insider trading’ restrictions that ban company officials and media representatives from dealing in securities for 12 hours after they learn earnings results, the Tokyo Stock Exchange wants companies to stop sending such results over the Internet right after they’re announced in news conferences.” Prior to the Internet, officials could make money on financial news. They had to wait 12 hours, but the news spread slowly, so they retained an advantage. The Internet gets the news out quickly; by the time officials can trade, it is too late.

The FBI and Tokyo exchange responded to potentially disruptive effects of efficient information dissemination by making efforts to stop or slow its spread. Such responses are unlikely to succeed.

5.4. Informal Interaction: Endangered or Irrepressible?

Toda (1991) outlines several roles played by informal spoken conversation, such as gossip or chatting. These include establishing, maintaining, and repairing status. Essential to the effectiveness of chatting is that it is ephemeral and restricted to a small audience, that it occurs in the here and now.

Teenagers asked why they like messaging sometimes respond that they like its informality – unlike email, you don’t have to worry about spelling and grammar. This is ironic to those who used email in the 1980’s, because email was often lauded in the same terms – unlike written or typed correspondence, it was informal, you didn’t have to worry about spelling or grammar. At the first CHI Conference, John Seely Brown argued seriously that a “spelling decorrector” should randomly insert errors in email to insure informality.

Email changed when the cost of memory plummeted, making it economic to save email after reading and sending it. Email had been ephemeral, confined to the ‘now.’ True, a recipient could forward email to someone on the same network before deleting it. Senders generally assumed that recipients would not forward (or print) gossip email, assumed that their remarks were restricted to the here and now, and violations were a surprise. Thus email had much the same characteristics and could serve much the same role as chatting.

But people saw value in being able to consult it later. As memory prices fell and email became archived more often, it gradually lost its informality. Spelling correctors were added. But it is a slow process; the Iran-Contra and DOJ-Microsoft cases exposed email sent with assumptions of ephemerality, email that in some cases arguably contained the posturing that Toda describes as a central facet of chatting. Email was no longer confined to the here and now. But a need for informal ephemeral interaction persisted, to serve the functions described by Toda. Online chatting and instant messaging meet this need.

But what next? “I wish I had a record of that great comment made in that chat session...” An aid to productivity? Smith et al. (2000) describe a prototype “threaded chat” application that archives chat, commenting that they were “fencing in the last frontier” of ephemeral interaction. In a study of the system, people sometimes went back to correct the spelling of earlier comments. By being archived for possible later viewing, it lost informality. Whether or not this was the reason why, users were not enthusiastic about threaded chat.

It is not clear that archived chat will become the norm, but even if it does, it seems unlikely that it will turn out to be the last frontier for ephemeral, here and now, informal conversation. Informality serves important roles in social functioning. We are likely find a way to work it into our lives, to work around context-awareness systems that would record it, and thereby ensure, perhaps, a limit to the totality of captured context.

6. CONCLUSION

In this essay, I noted that context-awareness is a much broader domain than awareness of identity and location. When we represent context digitally we represent only some of it, which can change the impression it makes. This can lead to an effort to represent more context. At the same time, putting information on a network changes it fundamentally by giving it a much broader potential audience. It shifts activity that was often local to a global stage, altering its nature and its uses. This contributes to a shift away from action in the “here and now” that marked much of humanity’s history to action that has consequences across space and time.

One can view this pessimistically, but the rapid embrace of digital technology clearly reveals our confidence that the benefits—efficiency, safety, accountability, access to resources—far outweigh the costs.

I confess to a qualm about this potential immortality of digital information: The loss of second chances, fresh starts. In static societies people's reputations may have faded slowly, but in today's complex and dynamic world, the capability to begin over is useful. We have been able to recreate ourselves when we go to university, change jobs, or move, distancing ourselves from others' expectations and assumptions about us. This capability is disappearing. The web site a teenager creates today may be accessed far in the future by potential employers, colleagues, spouses, children. No one will fear that their past will catch up with them because we won't leave our pasts behind. To me this is daunting. Starting over can be a way to grow, to experiment and optimize one's potential. Without this opportunity, people may have to be more honest, brave, and tolerant of the inconsistencies and behavioral shades of gray that context-aware platforms will capture.

NOTES

Acknowledgments. The author is grateful for conversations around these topics with many people, especially Steve Poltrock, and also to Anoop Gupta, A. J. Brush and the MSR Collaboration and Multimedia Group. He is particularly grateful to Masanao Toda for copies of chapters of his manuscript in progress.

Support. The author received support from National Science Foundation grant #IRI-9612355 for research cited in this essay.

Authors' Present Addresses. Jonathan Grudin, Microsoft Research, One Microsoft Way, Redmond, WA 98052 jgrudin@microsoft.com

HCI Editorial Record. (supplied by Editor)

REFERENCES

- Begeman, M., Cook, P., Ellis, C., Graf, M., Rein, G. & Smith, T. (1986): "Project Nick: Meetings augmentation and analysis", Proc. CSCW'96. Revised version in *Transactions on Office Information Systems*, 5, 2, 1987.
- Brin, D. (1998) *The transparent society: Will technology force us to choose between privacy and freedom?* Perseus.
- Dey, A. K., Abowd, G. D. & Salber, D. (2001). A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction*, 16, xxx-xxx.
- Erickson, T. & Kellogg, W. (2000). Social translucence: An approach to designing systems that support social processes. *Transactions on Computer-Human Interaction*, 7, 1, 59-83.
- Friedman, B. (1997). Social judgments and technological innovation: Adolescents' understanding of property, privacy, and electronic information. *Computers in Human Behavior*, 13, 3, 327-351.
- Isaacs, E.A., Morris, T. & Rodriguez, T.K. (1994). A Forum For Supporting Interactive Presentations to Distributed Audiences. *Proc. CSCW '94*, 405-416.
- Jancke, G., Venolia, G.D., Grudin, J., Cadiz, J.J. & Gupta, A. (2001). Linking public spaces: Technical and social issues. *Proc. CHI 2001*, 530-537.
- Johnson, J., Roberts, T.L., Verplank, W., Smith, D.C., Irby, C.H., Beard, M. & Mackey, K. (1989). The Xerox Star: A retrospective. *IEEE Computer*, 22, 9, 11-29.
- Lenat, D. B. (1995). Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38, 11, 32-38.
- Mark, G., Grudin, J., & Poltrock, S.E. (1999). Meeting at the desktop: An empirical study of virtually collocated teams. *Proc. ECSCW'99*, 159-178.
- Murphy, R. F. (1964). Social distance and the veil. *American Anthropologist*, 66, 1257-1274.
- Palen, L. (1999). Social, individual & technological issues for groupware calendar systems. *Proc. CHI 99*, 17-24.
- Smith, M., Cadiz, J.J. & Burkhalter, B. (2000). Conversation trees and threaded chats. *Proc. CSCW 2000*, 97-106.
- Teasley, S., Covi, L., Krishnan, M.S. & Olson, J.S. (2000). How does radical collocation help a team succeed? *Proc. CSCW 2000*, 339-346.

Toda, M. (1991). Common sense, emotion, and chatting, and their roles in interpersonal interactions. SCCS Technical Report 90-1-01. Chukyo University.

Toda, M. (1999). *The urge theory of emotion and social interaction*. Unpublished manuscript. Chukyo University.

Trigg, R., Suchman, L. & Halasz, F. (1986). Supporting collaboration in Notecards. *Proc. CSCW'86*.

Want, R., Hopper, A., Falcao, V. & Gibbons, J. (1992). The Active Badge locator system. *ACM TOIS*, 10, 1, 91-102.