Security in the Wild: User Strategies for Managing Security as an Everyday, Practical Problem

PAUL DOURISH[*], REBECCA E. GRINTER[+], JESSICA DELGADO DE LA FLOR[*] AND MELISSA JOSEPH[*]

**School of Information & Computer Science*
University of California, Irvine
Irvine, CA 92697-3425, USA
jpd@ics.uci.edu

[+]*Palo Alto Research Center (PARC)*
3333 Coyote Hill Road
Palo Alto, CA 94304, USA
beki@parc.com

Abstract: Ubiquitous and mobile technologies create new challenges for system security. Effective security solutions depend not only on the mathematical and technical properties of those solutions, but also on people's ability to understand them and use them as part of their work. As a step towards solving this problem, we have been examining how people experience security as a facet of their daily life, and how they routinely answer the question, "is this system secure enough for what I want to do?" We present a number of findings concerning the scope of security, attitudes towards security, and the social and organizational contexts within which security concerns arise, and point towards emerging technical solutions.

Introduction

Weiser's vision of ubiquitous computing – a third wave of computation, displacing the era of mainframes and personal computers – implies radical transformations in many aspects of our computational world (Weiser, 1991; 1993). By moving interaction beyond the desktop, it transforms the settings within which interaction occurs, and the forms of that interaction; by emphasizing the role of trends in miniaturization and power consumption, it transforms the nature of the computational devices themselves. At the same time, it also transforms the nature and boundaries of the "system." Where conventional computer use is focused on a single device, perhaps linked to a few others across a network, ubiquitous computing is typically manifest through collections of many devices – some mobile, some static, some embedded in the infrastructure, and some carried by individuals – brought together to form ad hoc coalitions in specific

circumstances of use (Edwards et al., 2002). Holding a meeting in an interactive workspace may involve bringing together tens of devices or more, including mobile, handheld and wearable devices belonging to meeting participants, as well as components managing the input, monitoring, recording, and display capabilities of the space (e.g. Johanson et al., 2002). Ubiquitous computing, then, implies ubiquitous digital communication, as the devices that make up a ubiquitous computing system communicate in order to identify each other and their capabilities, achieve effective configurations of functionality, and interoperate in support of user needs.

However, while ubiquitous communication offers the possibility of achieving more effective coordination in a world of computational devices, it also introduces a range of problems regarding the security of these systems. Information system security has always been an important issue in military and corporate settings, but in mobile and ubiquitous computing settings, it becomes a central concern for casual and end users. Networked and e-commerce systems bring with them the dangers of disclosing credit card numbers, social security information, bank transaction details, client records and other electronic artifacts; context-aware and mobile systems carry with them the possibility of disclosing information about activities and locations. Ubiquitous computing, as Weiser noted, anticipates that an individual's computational needs will be met by tens or hundreds of computational components working together; security is both an inherent problem in this sort of combinatorial system, and a practical concern for end users. Systems must be not only secure, but usably and practically secure.

In order to understand security as a user concern as well as a technical concern, our approach has been to look at the practical, everyday aspects of security as they manifest themselves in the use of current computer systems. We have been especially concerned with wired and wireless networked systems, which, although typically on a smaller scale, exhibit some of the same combinatorial problems to which ubiquitous computing systems are subject. By security, here, we refer to the ability of users to express and rely upon a set of guarantees that a system may make, explicitly or implicitly, about its treatment of user data and other resources. As far as possible, we distinguish between security, which is a largely technical concern, and privacy, which is a largely social concern; clearly, however, they are related, and we will return to this relationship towards the end.

Technical and Human Factors Approaches

The problem of usable security is not new to ubiquitous computing; the relationship between usability and security has been persistently troublesome. Security systems typically attempt to introduce barriers to action (such as passwords or other authentication mechanisms) while HCI designers attempt to remove such barriers. The barriers that security mechanisms erect often seem to confuse or frustrate end users, who then become the "weak links" in the security chain. Consequently, researchers have explored how to improve the usability of security systems, approaching this from both the bottom up (focused on security technology) and from the top down (focused on human factors.)

Most research on security in ubiquitous computing has taken the technical approach (e.g. Balfanz et al., 2002; Kindberg and Zhang, 2003; Stajano, 2002.) From the technical perspective, one approach has explored "transparent" security infrastructures – systems that are secure without the user needing even to be aware (Blaze, 1993). While this offers the potential advantage of freeing the end-user from the need to understand the mathematical concepts that form the basis of the security solutions they will use, it comes with drawbacks also. One difficulty with transparency is that does not help users "see" how to manipulate their security environment when the unexpected happens or when a computer system requires troubleshooting.

Another technical approach attempts to characterize different degrees of security provision, as embodied by the idea of "quality of security service" (Irvine and Levin, 2001; Spyropoulou et al., 2000). The fundamental insight is that organizations and applications need to trade-off different factors against each other, including security of various forms and degrees, in order to make effective use of available resources (Thomsen and Denz, 1997; Henning, 1999). While this work is directed towards resource management rather than user control, it begins to unpack the "security" black box and characterize degrees and qualities of security. Again, however, it requires users to understand the impacts of their actions and the workings of the security technology.

The second perspective has focused not on security technology but on *human factors* analysis. For example, Whitten and Tygar's (1999) usability analysis of PGP 5.0 demonstrated the difficulties that users have in completing experimental tasks; in their study, only 3 out of 12 test subjects successfully completed a

standard set of tasks using PGP to encrypt and decrypt email. This study highlighted the poor match between the design and implementation of the encryption technology and the users' understanding of how to secure electronic mail messages. In a series of studies, researchers at University College, London have explored some of the interactions between usability and security (Adams, Sasse and Lunt, 1997; Adams and Sasse, 1999). They have focused particularly on user-visible elements of security systems, such as passwords; the specific problems that they identify with passwords have also led to interesting design alternatives (Brostoff and Sasse, 2000; Dhamija and Perrig, 2000). Others (e.g. Yee (2002) or Zurko and Simon (1996)) have also explored the relationship between interaction design and security, and attempted to draw out a set of general interface design guidelines.

Explorations of technology and human factors highlight significant problems concerning achievable levels of security in electronic settings. However, we also know that, even when technologies are available and usable, they are still not always used, for a host of reasons. Accordingly, we feel that it is critical that we look not just at *technologies* but also at how those technologies are incorporated into contemporary security *practice*.

Some investigations of this sort have been carried out. For example, Freidman et al. (2002) have investigated users' perceptions of security in Web-based interactions, and note that even in this restricted domain, problems arise in how people assess the security of settings they encounter. Rimmer et al. (1999) and Sheeran et al (2001) discuss the mental models that end users develop of network structure and behavior, and illustrate the impacts that these have on system use. Weirich and Sasse (2001) present a preliminary investigation of users mental models of security; their investigation is similar to ours, although they focus largely on technologically sophisticated users.

Our research here builds on these investigations. We also believe that usable security is not just a matter of better interface design. Further, we proceed from the observation that security is an issue that people already deal with in their electronic workplaces, and so it is these current practices that we want to examine here. Through qualitative analysis, we are interested in uncovering how end users go about managing security as an *everyday, practical problem*.

Approach

We used two qualitative research approaches, semi-structured interviewing (Bernard, 1988) to gather the data and Grounded Theory (Glaser and Strauss, 1967) to analyze it. Grounded theory provides a set of procedures for developing analytic accounts of qualitative data, based on the iterative generation, validation, and refinement of coding schemes. A qualitative approach is more appropriate than a quantitative at this stage, given that our goal is not to provide definite answers to definite questions, but rather, to determine what questions we might want to ask in the first place. In particular, the goal of our investigations is not simply to document what users do, but rather to understand their experience of security as they encounter it. It is this concern – to be able to see security as it appears to the users we are studying – that drives our methodological approach. From this perspective, we gain more by understanding the experience of a small number of users in depth and detail than we would from a broader statistical account of the activities of a larger number of people. Our goal here is not to quantify user behavior, but to characterize and understand it.

While mobile and wireless network access did arise repeatedly in the course of our interviews, we did not focus solely on ubiquitous computing topics. There were two main reasons for this. First, we believe that the problems of security and information technology are broader than simply those encountered in ubiquitous computing, and we hoped to gain a more holistic understanding of these questions by looking at the full range of user activities. Second, we felt it inappropriate to impose our categorization of technology types onto user experience; while, as researchers, we might distinguish ubiquitous computing technologies from other, more conventional information systems, end-users do not adopt those categorizations of their own natural activity.

Sites

The results presented here have been collected from a series of qualitative investigations conducted at two sites over the past twelve months or so. The sites, and individuals, were selected to cover a range of types of work (from students at universities to corporate lawyers), and different backgrounds (familiarity with

computers or not), as well as more traditional types of difference (age and gender).

Site A is an academic institution. Our interview subjects here are administrative staff members of both an academic department and a research institute, and graduate students in a management program. We were interested in these people because of their range of institutional affiliations and responsibilities; previous research (Sheehan, 2002) suggests that new graduate students should be an interesting subject pool. We interviewed a total of eleven participants at Site A.

Site B is an industrial research lab. Here, we were particularly interested in interviewing people who had security needs because their jobs required certain levels of confidentiality resulting from institutional (Federal Government) and organizational (the corporate owner of the site) rules and policies. In other words, we focused on people whose data needed protection not just from the outside world, but from other people within Site B. This led to a focus on people in support and administrative work. At site B we conducted nine interviews.

In what follows, we discuss our findings from these sites in two clusters of related topics. The first concerns the ways in which security issues arise for people in the course of their regular activities, and how they understand them – we refer to this as the *experience of security*. The second is focused not just on interpretations but on actions, and what practices and patterns people adopt to manage their security needs and accommodate them into their work – we refer to this as the *practice of security*. Practice and experience cannot be entirely separated, but they provide a frame for beginning to tease apart relevant issues. In each section, we report the broad patterns we saw emerging in the data analysis, and use examples from our interviews to illustrate the findings in more detail.

The Experience of Security

Although the mathematical and technical foundations of security systems delimit the scope of "security" for the research community, end users see their encounters with security quite differently and set the scope of concerns more broadly.

Security as a Barrier

Our particular concern is, clearly, with security needs and security technologies, but we found that, quite consistently, respondents would expand the scope of

discussions to include other related issues, such as unsolicited email. While we initially found this frustrating, further analysis highlighted the ways in which these various items are metaphorically related and, as elements of everyday user experience, occur in combination. For our subjects, security and spam are two aspects of the same problem; as practical problems, viruses, network scanners, password sniffers, and unsolicited email form a "natural class," even though they may be technically quite different. What is more, conventional users not only regard these as the same problem, but also think of the same technologies as providing solutions; firewalls, for example, are described both as technologies to keep out unwelcome visitors but also unwelcome email messages. In other words, people seem to both imagine and seek unitary solutions to these problems.

When we think of the real-world experiences on which people base their experiences, they think of security as a barrier, akin to a gate or a locked door. Security is, generically, something to "keep things out," and so the various threats – the things that are being kept out – become co-constructed as the common entities against which security protects.

There are three immediate implications of this observation. The first is that a security solution that solves only one problem but not others (e.g. a virus scanner) could be potentially rejected by end-users for being only a "partial" solution to security problems *as experienced* by end-users. Conversely, a second observation is that a technology deployed to solve one problem may be mistakenly interpreted as providing protection against the others; for example, one user at Site A talked of being protected from viruses by a new filtering system installed by the network service (although, in fact, this was a spam filter with no virus detection facilities.) Third, the focus on barriers or "choke-points" diverts attention from channels, as exemplified, for instance, by users who install advanced firewalls but then run unencrypted 802.11b wireless networks. Taken together all of these reveal how the gap between theoretical security and effective security may be experienced by end-users in the form of expectation failures, mistaken assumptions, and a focus on one aspect of the problem blinding users to other security issues.

Online and Offline

The relationship between online and offline experience is a complex one, but is also centrally important. Online conduct seems to be continually shaped by

aspects of the offline world. We found two examples of the relationship between online and offline in our data.

One relationship is the potential leakage of information between online and offline settings. While our initial expectation was that people would relate Internet security problems to internet-based fraud (e.g. forging email, identity theft, unauthorized financial transactions), a much more immediate concern for a significant number of our subjects was the possibility that inadvertent information disclosure *online* could create a threat *offline*. Most frequently, these were problems of personal security. Stalkers were an especially common reported threat, and much of people's attention to online information disclosure concerned information that might result in direct personal threat. We were struck by the regularity with which this issue arose. Perhaps unsurprisingly, this observation was especially common among women.

The second type of relationship between online and offline aspects of security is marked by occasions where people must deal with the physical manifestation of their computer network. By computer network we mean the cables, routers, modems and other pieces of equipment through which their connection is maintained. Wireless networking technologies are perhaps especially interesting here due to their combination of tangible and intangible elements. Wireless networks offer people the same infrastructure interfaces that they conventionally associate with wired or point-to-point networks that conventionally carry with them obvious properties of physical security and accessibility. At the same time, however, they make the actual infrastructure of service provision intangible. One user we encountered spent some time trying to diagnose an apparent problem with a networked printer that refused to accept print jobs, before noticing that, in fact, he was connected through his neighbor's access point rather than his own (and so was connected from an unauthorized IP network). The very intangibility of network infrastructure makes it harder to for people to relate online experiences to offline manifestations of technology.

In other words, what we see from our observations is that, for everyday users, security is not purely an online matter; it extends into the physical world. The information which is to be protected, the resources to be managed, and the work to be carried out all exist in a physical setting too. Security practices may draw as much on the arrangement of physical spaces as on the arrangement of technical

resources and, again, providing people with technical solutions that cannot be understood or integrated into what people see as the "whole problem" will reduce their effectiveness.

Attitudes Towards Security

Clearly, experiences with security and systems vary between individuals. Our data suggest a range of attitudes that people display towards security.

Frustration. Although we did not specifically sample for age, our data suggest that age and experience are correlated with attitudes towards security. Specifically, our younger subjects, with a relatively longer exposure to computer systems (and in particular, it seems, childhood exposure) express a much greater confidence in their abilities with computer systems. In particular, they are more likely to report encountering situations in which security services proved problematic, hindering rather than helping their activities. Getting files through firewalls, for instance, had been problematic for some, who found that they had to turn off or circumvent security technologies in order to get their work done. They were more likely to talk of security in terms of its costs as well as its benefits, and frame technical security measures as ones that can interfere with the practical accomplishment of work. This is, of course, the "barrier" argument when seen from the other side.

A similar example occurs in a related study of teen use of SMS (text messaging via GSM's Short Message Service) reported elsewhere (Grinter and Eldridge, 2003.) The teens studied never intentionally turned off their phones, which meant that they rarely if ever used their password to log back onto the phone after a reboot. This meant that when they accidentally let the battery run out, the teenagers would need to take their mobile phone to the nearest service center to get the password reset before they could resume exchanging SMS's. This delay, in addition to creating inconveniences to have to go to the service center via public transportation, also caused consideration frustration when the teenagers realized just how many messages and potential activities they would miss out on in the time it would take to get fixed.

In other cases, security appears as an obstacle in other ways. For much day-to-day use, security is not a primary concern for end users; people rarely boot their computer in order to deal with security configurations. The persistence of virus

checkers, intrusion detectors, and other similar systems in interrupting current work in order to insist that something be done (new rules installed, ports blocked, or even just a button pressed) seemed to be problematic. This is, perhaps, another case of the difficulty of an "all-or-nothing" approach – security is either something unmentioned, or it is something to be dealt with suddenly and immediately.

Pragmatism. In broad terms, and in line with the previous observation, the younger respondents seemed more pragmatic about their security needs, expressing more nuance about the situations in which they might need security. For instance, they discussed using known insecure technologies in settings where they felt that the risks were justified (e.g. a machine that was known to be chock full of viruses, but was otherwise unused so it didn't matter.) This pragmatic orientation in younger subjects is in line with previous findings (Sheehan, 2002). Pragmatic users see security as a trade-off, one that must be continually struck as one balances immediate needs against potential dangers. For pragmatic users, then, systems need to be both flexible and translucent, so that these trade-offs can be made effectively.

Futility. However, even amongst those who expressed more confidence about their abilities and a more pragmatic orientation towards security, there is an overwhelming sense of futility in people's encounters with technology. This corroborates the similar observation was made by Weirich and Sasse (2001) in their investigations. Our subjects make repeated reference to the unknown others (hackers, stalkers, etc.) who will "always be one step ahead," and whose skill with technologies will mean that there are always new attacks to be diverted. As a result, they talk repeatedly of security lying not so much in technology as in vigilance; the continual, active defense against new and evolving threats.

The results of this sense of futility vary depending on the setting and the forms of threat. With respect to the broad Internet, it certainly contributes to frustration and the sense that one is continually "running to stay in the same place"; it creates a fictive norm of adequate protection, against which people continually find themselves wanting. In organizational settings, it becomes manifest mainly as a concern with "due diligence" – the visible demonstration that one has done enough. As in the cases discussed earlier where security moves out of the computer and into the physical environment, it becomes important to demonstrate

to others that one has taken due care to manage information and activities securely, even though respondents may not feel that these measures are likely to survive an assault.

The Practice of Security

In the previous section, we illustrated the conflicting interpretations of security as it presents itself in users' everyday encounters with technology – as a protected barrier, as a problematic impediment, as necessary, as futile, etc. However, despite the various interpretations, security cannot be ignored, but must be dealt with; a concern for security, or a means for dealing with it, must be integrated into how people deal with technologies. In this section, we consider the practices of security – how concerns for security are integrated into the details of how people conduct their online activities.

Delegating Security

Unsurprisingly, most people, in the course of their daily work, have neither the time nor inclination to be continually vigilant for new threats; they are focused on getting their work done. One particularly interesting issue, then, is the various modalities by which people delegate responsibility for security. Security is, to some extent, turned into someone else's problem, or at least, external resources are marshaled to deal with the problem. Four forms of delegation are identifiable in our interviews.

The first is to *delegate to technology*, which involves relying on some form of technology for protection. So, people might rely on SSL encryption for data connections, use ssh tunneling for their email, or trust that switched Ethernet is more secure than a traditional common medium. These are, of course, the solutions that the technical community is used to providing. Interestingly, though, this was perhaps one of the least common ways of managing security that we encountered. Perhaps this was because the only people who tended to be confident enough to make that argument in our interviews were those who understood the security vulnerabilities that these technologies prevented, and what the limitations of each was, and how to use them in combination to create secure solutions. In other words, it was an argument that was only available to those who could in fact turn a technically working security system into an individually workable solution.

It is also interesting to observe that this delegation is an investment of trust, and we speculate that it depends on visible presence of technology to be trusted, which questions the idea of security as an invisible or transparent facet of a system. Arguably, the use of physical arrangements to embody security concerns is a case of delegation to technology (albeit less advanced.)

The second mode of delegation is to *delegate to another individual*, such as a knowledgeable colleague, family member, or roommate. This tended to emerge in interviews where people were describing personally owned devices. The knowledgeable person also tended to be the person who had helped them in a previous context, such as in discussing what to get, helping them set up the computer. The knowledge and skill of the "technical friend" was cited as one element of a person's defense against potential threats. For people who feel limited in their ability to assess technology, known individuals may be more trustworthy, often grounded in a series of positive experiences.

The third mode is to *delegate to an organization*; like delegation to an individual, this delegates to others, but the others are organizationally defined and may not even be known personally. The most common form of this was the "we have a very good support group" argument. More generally, it is the skills and especially the vigilance of the organization in which people place their trust. In some cases, again due to the fictive norm associated with vigilance, more trust may be accorded to external organizations; and so, facilities run by a central service rather than by a local group, or facilities managed through an outsourcing arrangement, are seen as more secure.

Finally, we also found a mode in which people would *delegate to institutions*. Like Friedman et al. (2002), we also found that people would tend to trust that certain types of institutions, such as financial institutions, would take appropriate security measures. Again, this is an online/offline relationship; impressions formed about institutions such as banks' concern with physical security (locked vaults and armed security guards) are carried over to online security, even though of course online interactions with a bank depend on a complex of intermediate technologies outside of any bank's control.

There is an important temporal aspect to this process of delegation. Essentially, once responsibility has been delegated, it becomes almost invisible; it seemed to be rare for these issues to be revisited. Individuals to whom responsibility had

been delegated when they set up the computer sometimes disappeared from view (the former roommate or colleague, or the son who had left for college), and yet they were still invoked as the guarantor of security. In organizational settings, we found that, over time, some newer employees would not know what kinds of access controls, for example, would be on their file systems. Delegation to the support group would have occurred several years prior to their arrival and they could not articulate what kinds of privileges existed. This is interesting for two reasons. First, the work practices of groups often "grow over" the underlying security while taking advantage of what it provides, until the security decisions are lost to conscious memory. Second, no-one concerned themselves with this. Between the initial security decision and the supporting work practices, the day-to-day configuration had disappeared but was still being enacted correctly.

Secure Actions

The people we interviewed had a number of methods for managing online security, some of which involved using security protocols in what may seem like unusual ways, and others of which that appear to involve no security, but illustrate how people think about security in technology.

Whitten and Tygar's (1999) analysis of email discovered that users experience considerable difficulty using PGP to secure their communications. However, what is clear is that people use email to communicate all the time, and even when they have information that needs to be protected. So, we wondered what they did to "secure" the information. We discovered two common strategies.

First, people use institutional means to secure communications and our informants were no exception to this practice. We see this each time we receive an email from someone that has an attached signature file that states the legal and illegal uses of the contents of the message. Rather than securing the communications, the purpose of these statements is to defend the contents if they become incriminated. In other words, corporations often attempt to mitigate the risks of information leaks by securing the consequences of those leaks by marking the messages. Although it may not be secure technically, it is not surprising that this approach is used. Marking documents has long been a means by which corporations sort and prioritize the contents of presentations and reports and so forth. The securing of

email messages appears to coincide with the migration of email from an informal chatting technology to a formal means of corporate communications.

Second, we also found cases where people were using context to secure (or perhaps "obscure") their email messages. By this we mean that we found cases where people described sending email that did not explicitly state what the subject was in the actual email itself, but used a shared working context to express the new information. For example, an email message that says, "I took the actions you requested" could refer to many types of activity. Moreover, by removing any sense of time from the contents (other than the date stamp) no specific temporal information could be deduced.

Crucially, this arose not simply as happenstance; rather, it was an explicit strategy adopted to support secure communication as a part of a broader pattern of work. Using "cryptic" email was simply a lot easier to do than using a security tool to encrypt the information. By easier though, we do not just mean Whitten and Tygar's usability of various encryption software, we mean that context-based encryption may simply be the more visible form of security measures for many people working with email. The fact that it can be accomplished as part and parcel of the working activities, rather than as a separate and parallel activity, also makes it a more convenient feature of work practice (Smetters and Grinter, 2002).

The visibility of security systems (their presence and utility) let alone their usability is also illustrated by the use of media switching as a security measure. In our interviews, we found several occasions where people switched communications media based on security decisions. Several interviewees said that they trusted the telephone for their most secure conversations, and introduced a media switch to the telephone from email when the most sensitive of topics came up.

Perhaps what is most surprising about this ability to rate technological mediums for security is that the same phenomenon is reported in other settings (Grinter and Palen, 2002). Teenagers using Instant Messaging technologies also reported suggesting a media switch to the telephone for the most confidential of conversations. The telephone offers two related advantages. First, potentially, it is a less vulnerable medium than email. Although it can be tapped and listened into, maybe it is less statistically likely. Second, and as more commonly articulated, the medium is ephemeral in the sense that nothing that is said is readily recorded and

transmittable. Unlike electronic text, something that teenagers observed with clarity was that it was much harder to record and convince anyone else with absolute certainty what was said on the telephone. In other words, confidentiality and privacy of information can be more likely guaranteed in a medium that is not as readily recorded, and understanding the different between electronic conversation and electronic text, the audio word seemed more secure than the textual one.

Earlier, we discussed encryption and the need for secure communications and its relationship to media choice. In that and previous discussions, current security systems seemed almost to fail our users in the sense that they did not fit work practices. Moreover, they competed ineffectively with other alternatives such as simply switching media (something that was once just available to office workers is now something that teenagers at home can consider). However, we also found some occasions where security measures had been incorporated into working practices.

One example of this concerns access control to shared directories. In this case, two legal staff explained how they used the access control settings for online directories as a means of communications. Specifically, they both worked on files that once they had finished with their own notes needed to be sent to a centralized legal body for further processing. Rather than using email to do this, they used a system of shared online directories. While they worked together and locally on the files, they put them in an online directory that only they could access. When they had finished working on the file they would simply move the file to another directory, one whose access controls were set to allow other people to access it from a remote site. One advantage that the two local legal staff found with this scheme was that they did not have to know specifically to whom they had to send the files (unlike email).

Holistic Security Management

Another strategy using for managing online security relies upon the physical arrangement of space. The richest example of the use of space to protect online data came from a person whose job required her to manipulate confidential data and simultaneously frequently interact with individuals who came to her office. As a consequence of her work, she had spent considerable time thinking about the

relationship between online and physical security. For example, she positioned her computer screen to face her, and critically, to point away from the first point of entry into her office. If someone showed up to turn in a document or start a conversation they could not see potentially sensitive information on her screen. Moreover, her online work often required her to refer to physical documents, which could also contain sensitive information[1]. She simultaneously needed to keep documents to hand, by the monitor, but not allow them to be seen by others. Consequently, she had devised a system of colored folders that allowed her to store documents out of sight, but did not require labeling (since that alone would sometimes be enough to cause a breach in security). She used the different colors to represent different actions required. The colored folders allowed her to balance security (by preventing people from seeing the documents or identifying types of activity) and information access (allowing her to see the state of work and know where various materials were stored).

Finally, her practices were all supported by her office layout, which she used as another means of protecting information. She had arranged her office (which she alone occupied) into two distinct sections. On entering the office, the first section was where office visitors would sit. All the spare seats were in this part of the office, and there was a barrier (in the form of a desk) between the public visiting part of the office and the second – her private – part of the office. In the front part of her office, the surfaces were clear.

By contrast, the back part of the office contained a number of file folders. Although, in theory, someone who wanted to access the back of the office could do so, social conventions prevented people from getting up from the guest seats and walking around the desk into the back part of the office, and were used to regulate visitors' access to certain parts of the office and the data stored there. The nature of this person's work made her very articulate about the relationship between online and offline security. This interview lead us to follow these lines of questioning in other interviews. We saw similar patterns in other people's physical arrangements. For example, we observed that several administrative assistants to executives who routinely read and managed their bosses' email kept their monitors turned away from the casual gaze of visitors. We also saw

[1] It should be noted here that these were the practices she used for managing outstanding work items, physical documents that had been finished with were stored under lock and key.

managers' offices that had clearly demarked zones for visitors and meetings, and separated areas for private work where, again, the computer typically resided. The use of physical space suggests that people do not think about online security in isolated terms, not just the potential vulnerabilities of the data they manipulate at their terminals, but also in the defensive mechanisms they take to protect that data. However, beyond the manipulation of furniture, these spatial arrangements capitalize on important social conventions that determine what constitutes acceptable behavior as played out over a small, but significantly demarked, space. The manipulation of these conventions in this space, suggests another type of security expertise, one associated with the design of holistic security solutions for everyday practice.

Managing Identity

In the interviews we discovered another challenge for security: that of identity. This manifested itself in two ways – the production of identity, and the interpretation of identity.

First, we found that our informants were very conscious of the ways in which they presented themselves online. Many, for instance, maintain many virtual identities (email addresses, online personas, etc) as a way of controlling their visibility. In addition, we found some evidence for the use of *partial* identities; by controlling which specific aspects of information (social security number, home zip code, etc) they gave to different entities, some respondents attempted to maintain control over the degree to which they could be identified and tracked. When identity is seen as the sum of these factors, a subset seems secure.

The second issue is the interpretation of identity. In many proposed solutions an individual manages their own security. Personal firewalls, personal encryption, passwords, and so forth all place a primacy on the individual. However, we found a number of cases where seemingly individual people were in fact groups of people, and because of that security decisions and solutions became more problematic.

This problem was most clearly exhibited by individuals who had personal assistants. In many cases, the email address of an executive does not equate to the person themselves: it refers to them and their personal assistant. When we talked with assistants and people who emailed executives we discovered that

mismatched expectations presented difficulties with this arrangement. Although turning an executive's email into a group distribution list (that goes to the executive and their assistant) we found cases where people were surprised by this distribution. This surprise typically arose, because often email handles help individuals determine whether an address belongs to an individual or a group. For example *bob@company.com* is expected to go to a person by that name alone, but *soccer@company.com* could be considered as a group email for people interested in football.

Apparently individual entities which turn out to be collectives, and apparently collective entities which turn out to be individuals, place pressures on the mechanisms that allow people to control and manage information disclosure. Inversely, as Palen and Dourish (2003) observe, people act continually and simultaneously in multiple capacities – as individuals, as representatives of organizations or professional groups, as family members or members of some occupational groups, etc. The ways in which they act – the information that they chose to disclose, and how, when, and to whom they disclose it – mark them as affiliated with one or another set of people. Conventional separation into "roles" fails to capture the fluid and especially the simultaneous nature of these capacities in which one acts.

Reframing Security for Ubiquitous Computing

As we noted at the outset, security has long been noted as a significant problem for ubiquitous computing. Inherently, the distribution of interactive services across a range of devices and computational elements in different ownership and administrative domains implies a much more robust approach to security than is normally adopted in conventional software systems. At the same time, many ubiquitous computing applications involve personal data of one sort or another (such as location information), which must be protected.

Our focus on security as a practical problem suggests an alternative way of framing the security question in the first place. Traditional technological approaches regard the fundamental security question as, "what sorts of mathematical and technical guarantees can be made about the interaction between these components and channels?" By contrast, our investigations suggest approaching security from the other side, and thinking of the fundamental security

question as a user question – "is this computer system secure enough for what I want to do now?" This is a question that people routinely encounter and answer (if not always correctly or accurately); the problem for ubiquitous computing is to ensure that people have the resources they need to adequately answer the question.

In the course of presenting our empirical findings above, we have noted a number of specific design considerations that arise. In this section, however, we want to step back to consider broader design concerns.

We have tried to show here that security arises for many people as a primarily practical concern. For end users, getting their work done with a variety of technologies at hand, security is simply one of a range of considerations that impinge upon the practical accomplishment of work. Decisions about security, then, arise *in the course of* specific sequences of action, and *in the context of* a range of physical, social, organizational, and practical considerations. Indeed, it may be *inherently* implausible for typical users to specify, in advance of particular circumstances, what their security needs might be; those needs arise only as a result of specific encounters between people, information, and activities.

However, most conventional approaches to system security – based on access control lists, firewall configurations, host- or network-based access policies, etc – separate security-related decision-making from working activities that those security policies are intended to support. A primary design goal for ubiquitous computing technologies, then, is *to be able to place security decision-making back within the context in which it makes sense as a practical matter*, and so to design for security as an aspect of everyday activity.

In Ubicomp, this should not come as a surprise. One of the premises upon which ubiquitous computing is founded is that traditional desktop computing technologies decontextualize working activities; ubiquitous computing seeks to use emerging technologies to place computational activities back within the context of human action. So, context-aware computing, for example, represents an attempt to recontextualize computational activity by making it sensitive to the settings in which it takes place (Moran and Dourish, 2001.)

A broader set of design considerations, then, address the relationship between control over security and the activities being carried out with computer systems. Interestingly, while security technologies are often seen as ways of helping to manage people's privacy, we have found it useful to invert this relationship, and

use research into privacy management as an enlightening perspective in thinking about security.[2]

In particular, research into privacy management in electronic settings has noted the dialectic nature of privacy management (Altman, 1975; Palen and Dourish, 2003; Ackerman, 2000), and we can clearly see security emerging in a similar light in our data. Just as "privacy" is a dialectic process of boundary regulation, so the very definition of what counts as "secure" is a contingent matter; "security" depends on the circumstances. Information protection and information sharing mutually define each other, and that the process of managing security is the dynamic process of managing that balance. This perspective yields three high-level implications for design.

First, it suggests that protection and sharing of information are *two aspects of the same task*; they are always carried out together. For example, switching media from email to the telephone during a discussion or using cryptic email are simultaneously sharing and protecting information. Interestingly, though, most systems separate information protection and sharing whether it be for a single individual or among a group. Typically, information sharing is a primary task, while information protection is a subsidiary task, specified in advance through control panels, configuration controls, abstract rule specification, etc. Our investigations suggest that we need to think of these as being the same tasks; I should use the same mechanisms to share information as to protect it, and, critically, they should be available to me at the same time. A split in which sharing information (adding files to a server, sending an attachment, logging into an IM service) is separated from the work of protecting it (specifying access control, encrypting the information, specifying availability) is ineffective. Essentially, practices such as maintaining multiple email addresses or AIM screen names is a practical solution that users have forged which allows them to conjoin information sharing and information protection as a unified task.

Second, a critical issue that arises out of many of our observations is the extent to which people are able to monitor and understand the potential consequences of their actions. This was clear in cases where people described turning monitors away from the gaze of others. It also came out in cases where people did not clearly separate or grasp the subtle distinctions in security solutions (such as

[2] We are grateful to Mark Ackerman for pointing out this dualism.

installing a firewall and then running an unencrypted wireless network). Visibility of system behavior *on their terms*, or the lack of it, was often a reason that people understood whether something was secure or failed to realize whether something was not protected.

Since security requirements depend on the specific circumstances of action and are subject to continual reflection and revision, it is necessary to provide people with the means to understand the security implications of the current configuration of technologies at their disposal. Rather than being "transparent," then, security technologies need to be *highly visible* – available for inspection and examination seamlessly as a part of work. Moreover, that visibility needs to be expressed not as mathematically-grounded cryptographic concepts, but in terms that fit users' activities and needs at the time. Interactive system design emphasizes that available functionality and courses of action should be continually available at-a-glance to users; security configuration should be available in just the same way. This, critically, is quite different from being able to "call up" security information when it's needed; the point is that it is needed as part and parcel of every activity in the system (Dourish and Redmiles, 2002).

Third, and taking this one step further, we note that security is a *mutual achievement of multiple parties*. Like the people sending cryptic email to maintain the security of their information, people achieve security in the context of the activities that they carry out together. Even the colored folders were a mutual achievement by the person who designed the system and the other people she interacted with (who both provided some of that confidential data and needed to be prevented from seeing others' similarly sensitive information).

More generally put, the security consequences of my actions depend not only on what I do but also on what others' I interact with do. The scope of security, then, is a collaborative scope; it extends beyond the individual. Recognizing that the correct unit of analysis for security systems is groups rather than individuals, and that visualization approaches such as those suggested above might apply to groups, significantly changes how we conventionally think of security systems and security interfaces. Security is a collective accomplishment, an outcome of shared practices as well as shared configuration and technology; security management may be better conceptualized as a group task rather than an individual one.

Conclusions

Security and privacy are a "penultimate slide" problem for ubiquitous computing – the second to last slide in many conference presentations notes that there are security or privacy problems with the system that has been described, and that future work will address these issues. "Penultimate slide" syndrome indicates two things – first, that security and privacy issues are centrally implicated in the future success of ubiquitous computing, and, second, that they are currently poorly understood.

One feature of security management in a world of personal and ubiquitous computing is that it is, essentially, an end-user problem. Ubiquitous computing applications often depend on the spontaneous, ad hoc assembly of system components only some of which are under an end user's control, but which must be understood and arranged by end users to meet their needs. Effective security, in these contexts, must similarly be manageable and understandable to end users. While the research community has been extremely successful in developing the mathematical and technical foundations of secure computing and communication services, we have, perhaps, been less successful in the more practical task of making systems effectively secure. Any technology for secure communication is only as secure as the settings within which it is deployed. We have argued that a major obstacle to the development of more effective security strategies is that these systems often match poorly to the ways in which people need to make use of them. Accordingly, we have been attempting to uncover and understand the strategies by which users tackle system security as a practical, everyday problem. This investigation has begun to show a series of important considerations for design, and in particular has highlighted the importance of security management as a dynamic and intrinsic element of other activities.

A small but growing group of researchers have begun to examine the usability of security technologies, and have noted a range of problems that interfere with the effective use of technologies currently employed for security in day-to-day settings. However, we believe that effective solutions will not come solely from repairing the usability problems associated with existing technologies, because the very nature of those technologies – the ways in which they conceive of the problems of security – is a source of trouble. When we look at those areas in which HCI can be said to have lead to radical improvements in usability – such as

the development of graphical user interfaces and the development of the Internet into the Web – it is instructive to note that they did not arise through the incremental modification of existing technologies (e.g. systematically improving the usability of each command-line UNIX program.) Similarly, we suggest that effective security will require that we examine the conceptual models on which our systems are built.

As we have noted, the problems of security and usability are by no means isolated to the domain of mobile or ubiquitous computing. However, the specific issues that arise in ubiquitous computing contexts – pervasive communication, many devices, ad hoc configurations of technology, and "invisible" wireless communication infrastructures – make ubiquitous computing a particularly valuable test bed for experimentation. It seems quite likely that research into security in ubiquitous computing may have valuable implications for more traditional technology domains as well. Our investigations bear out the intuition that these problems are pervasive and pressing.

Acknowledgements

We would like to thank Mark Ackerman, Tom Berson, Brinda Dalal, Leysia Palen, David Redmiles, and Diana Smetters for their contributions to this research and this paper. We also gratefully acknowledge the patience and help of our interview subjects. This work has been supported in part by National Science Foundation awards IIS-0133749, IIS-0205724 and IIS-0326105, and by a grant from Intel Corp.

References

Ackerman, M. S. 2000. The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility. *Human Computer Interaction*, 15(2-3), 179-203.

Adams, A. and Sasse, M.A. 1999. Users Are Not The Enemy: Why users compromise security mechanisms and how to take remedial measures. *Comm. ACM*, 42(12), 40-46.

Adams, A., Sasse, M.A., and Lunt, P. 1997. Making Passwords Secure and Usable. In Thimbleby, H. O'Connaill, B., and Thomas, P. (eds), *People and Computers XII: Proceedings of HCI'97*, 1-19. Springer.

Altman, I. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory and Crowding*. Monterey, CA: Brooks/Cole Publishing Co. Inc.

Balfanz, D., Smetters, D., Stewart, P., and Wong, H. 2002. Talking to Strangers: Authentication in Ad-Hoc Wireless Networks. Proc. Network and Distributed System Security Symposium NDSS'02 (San Diego, CA.)

Bernard, H. R. (1988). *Research Methods in Cultural Anthropology*. Newbury Park, CA: Sage.

Blaze, M. 1993. A Cryptographic File System for UNIX. Proc. ACM Conf. Computer and Communications Security, 9-16. New York: ACM.

Brostoff, S. and Sasse, M.A. 2000. Are Passfaces more usable than passwords? A field trial investigation. In S. McDonald, Y. Waern & G. Cockton (Eds.): *People and Computers XIV - Usability or Else! Proceedings of HCI 2000*, 405-424. Springer.

Dhamija, R. and Perrig, A. 2000. Deja Vu: A User Study. Using Images for Authentication. In *Proceedings of the 9th USENIX Security Symposium*, Denver, Colorado.

Dourish, P., Edwards, W.K., LaMarca, A., Lamping, J., Petersen, K., Salisbury, M., Terry, D. and Thornton, J. 2000. Extending Document Management Systems with User-Specific Active Properties. *ACM Transactions on Information Systems*, 18(2), 140-170.

Dourish, P. and Redmiles, D. 2002. An Approach to Usable Security Based on Event Monitoring and Visualization. *Proceedings of the ACM New Security Paradigms Workshop 2002* (Virginia Beach, VA). New York: ACM

Friedman, B., Hurley, D., Howe, D., Felten, E., and Nissenbaum, H. 2002. Users' Conceptions of Web Security: A Comparitive Study. Short paper presented at ACM Conf. Human Factors in Computing Systems CHI 2002 (Minneapolis, MN.)

Edwards, W.K., Newman, M.W., Sedivy, J.Z, Smith, T.F., and Izadi, S. 2002. Challenge: Recombinant Computing and the Speakeasy Approach. Proc. Eighth ACM International Conf. On Mobile Computing and Networking (MobiCom 2002). (Atlanta, GA). New York: ACM.

Glaser, B. and Strauss, A. 1967. *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Chicago: Aldine.

Grinter. R. and Palen, L. 2002. Instant Messaging in Teen Life. *Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 2002* (New Orleans, LA), 21-30. New York: ACM.

Grinter, R. and Eldridge, M. 2003. Wan2tlk? Everyday Text Messaging. *Proc. ACM Conf. Human Factors in Computing Systems CHI 2003* (Ft Lauderdale, FL). New York: ACM.

Henning, R. 2000. Security Service Level Agreements: Quantifiable Security for the Enterprise? *Proc. New Security Paradigm Workshop* (Ontario, Canada), 54-60. ACM.

Irvine, C. and Levin, T. 2001. Quality of Security Service. *Proc. ACM New Security Paradigms Workshop*, 91-99.

Johanson, B., Fox, A., and Winograd, T. 2002. The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms. IEEE Pervasive Computing, 1(2), 67-75.

Kindberg, T. and Zhang, K. 2003. Secure Spontaneous Device Association. Proc. Ubiquitous Computing Ubicomp 2003 (Seattle, WA.) Lecture Notes in Computer Science LNCS 2864, Springer.

Moran, T. and Dourish, P. (eds.) 2001. *Context-Aware Computing: A special issue of Human Computer Interaction*, 16(2-4).

Palen, L. and Dourish, P. 2003. Unpacking "Privacy" for a Networked World. *Proc. ACM Conf. Human Factors in Computing Systems CHI 2003* (Ft. Lauderdale, FL). New York: ACM.

Rimmer, J., Wakeman, I., Sheeran, L., and Sasse, M.A. 1999. Examining Users' Repertoire of Internet Applications. In Sasse and Johnson (eds), *Human-Computer Interaction: Proceedings of Interact'99*.

Sheehan, K. 2002. Towards a Typology of Internet Users and Online Privacy Concerns. *The Information Society*, 18, 21-32.

Sheeran, L, Sasse, A., Rimmer J., and Wakeman, I. 2001. How Web Browsers Shape Users' Understanding of Networks. *The Electronic Library*, 20 (1), 35-42.

Smetters, D. and Grinter, R. 2002. Moving from the Design of Usable Security Technologies to the Design of Useful Secure Applications. *Proc. ACM New Security Paradigms Workshop NSPW 2002* (Virginia Beach, VA). New York: ACM.

Spyropoulou, E., Levin, T., and Irvine, C. 2000. Calculating Costs for Quality of Security Service. *Proc. 16th Computer Security Applications Conference*. IEEE.

Stajano, F. 2002. *Security for Ubiquitous Computing*. Wiley.

Thomsen, D. and Denz, M. 1997. Incremental Assurance for Multilevel Applications. *Proc. 13th Annual Computer Security Applications Conference*. IEEE.

Weirich, D. and Sasse, M.A. 2001. Pretty Good Persuasion: A first step towards effective password security for the Real World. *Proceedings of the New Security Paradigms Workshop 2001* (Sept. 10-13, Cloudcroft, NM), 137-143. ACM Press.

Weiser, M. 1991. The Computer for the 21st Century. Scientific American, 265(3), 94-104.

Weiser, M. 1993. Some Computer Science Issues in Ubiquitous Computing. Communications of the ACM, 36(7), 74-83.

Whitten, A. and Tygar, J.D. 1999. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. *Proc. Ninth USENIX Security Symposium*.

Yee, K.-P. 2002. User Interaction Design for Secure Systems. *Proc. 4th International Conf. Information and Communications Security* (Singapore).

Zurko, M.E. and Simon, R. 1996. User-Centered Security. *Proc. New Security Paradigms Workshop*. ACM.