

Lessons Learned From the Deployment of a Smartphone-Based Access-Control System

Lujo Bauer Lorrie Faith Cranor Michael K. Reiter Kami Vaniea
Carnegie Mellon University
{lbauer, lorrie, reiter, kami}@cmu.edu

ABSTRACT

Grey is a smartphone-based system by which a user can exercise her authority to gain access to rooms in our university building, and by which she can delegate that authority to other users. We present findings from a trial of Grey, with emphasis on how common usability principles manifest themselves in a smartphone-based security application. In particular, we demonstrate aspects of the system that gave rise to failures, misunderstandings, misperceptions, and unintended uses; network effects and new flexibility enabled by Grey; and the implications of these for user behavior. We argue that the manner in which usability principles emerged in the context of Grey can inform the design of other such applications.

Categories and Subject Descriptors

H.1.2 [Models and Principles]: User / Machine systems—*Human Factors*; H.5.2 [Information Interfaces and Presentation]: User Interfaces; H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces; K.6.5 [Management of Computing and Information Systems]: Security and Protection—*Authentication*

General Terms

Human Factors, Security

Keywords

Access control, usability, security, mobile computing, smartphones

1. INTRODUCTION

People use a variety of *access control* mechanisms to restrict access to physical spaces (e.g., rooms, cars) or electronic resources (e.g., computer accounts, web sites). Numerous access-control mechanisms have been developed: physical keys, proximity cards, and swipe cards are examples common for physical resources; passwords, RSA SecureID tokens,¹ and smart cards are more common for electronic resources. Some access-control technologies are used in both physical and electronic domains. Small electronic devices known as “fobs,” are used for access control to both computers and automobiles. Magnetic-stripe cards can serve as swipe

¹<http://www.rsasecurity.com/node.asp?id=1156>

cards for physical access or as credit/debit cards, conveying the authority to incur debt or perform a withdrawal. However, even when access-control technologies are theoretically capable of multiple uses, deployments are rarely interoperable. Thus, it is common for people to regularly carry a ring of physical keys, multiple magnetic-stripe cards, and one or more fobs, all while remembering multiple passwords.

Although some access-control mechanisms are quite familiar and most people are fairly adept at using them, they suffer from a number of drawbacks. For example, because keys must be physically copied and transferred, they are inconvenient for granting temporary access. One must plan ahead to grant access to others, making arrangements for keys to be copied and distributed prior to an access need. It is also nearly impossible to prevent unauthorized copying of keys, thus making it difficult to ensure that people granted temporary access to a resource can no longer access it after returning their keys. New electronic discretionary access-control technologies have the potential to address these and other drawbacks of familiar access-control mechanisms. However, these new technologies may introduce new problems that may undermine their usefulness.

In this paper we present an analysis of a trial deployment of a technology called Grey [4] that is designed to replace existing access-control technologies in a range of domains. Grey utilizes an off-the-shelf smartphone (augmented with the Grey software) as the user’s device for exercising her authority, and in our present implementation enables both computer logins (for Windows and Linux computers) and access to offices in our university building. Authority in Grey is represented by *credentials* held on the user’s phone, which the phone can present to a resource to gain access. In addition, Grey enables users to *delegate* authority to other users by creating new credentials and transmitting them via the cellular phone network.

In other work we compare the access-control policies implemented by Grey and keys, and quantitatively show that Grey allowed users in our field trial to implement their ideal policies more accurately and securely than they could with keys [3]. Here we focus on when and how users make use of the Grey system, and the obstacles they encountered.

We incrementally deployed Grey and recruited participants in an office building at Carnegie Mellon University. After nine months, our study encompassed 19 participants with Grey-enabled phones. We periodically interviewed the participants and tracked their usage of Grey via logs generated by the system. We report the primary results of our study as demonstrations of certain principles in our trial. The principles pertain to usability downfalls (failures, misunderstandings, misperceptions), network effects, and new flexibility offered via the Grey application, as well as the implications of the-



Figure 1: Left: Main screen containing all the resources the phone knows about, clicking on a resource causes the phone to attempt to unlock it. Right: A user proactively creating a delegation.

se for user behavior. While many of the principles themselves are generally held beliefs, how they emerged in our trial was in many cases unanticipated and illuminating. We thus believe these serve as useful lessons for those contemplating the deployment of a mobile application, particularly one on which users will depend for both security and access.

Our attention to smartphone-based applications is motivated by trends showing smartphones sustaining healthy market growth, e.g., a 75% increase in shipments worldwide from mid-2005 to mid-2006 [8]. Poised to inherit the existing cellular phone market, which has already reached vast worldwide penetration,² smartphones are likely to become the world’s first truly ubiquitous computing device. We believe that the lessons learned from our trial elucidate some of the challenges facing the deployment of advanced applications on this platform, and thus can expedite the design of such applications for a broad user population.

2. BACKGROUND

In this section we provide an overview of the Grey system and we discuss several attempts to evaluate the usability of access-control mechanisms.

2.1 Grey

Grey is a distributed access-control system that uses off-the-shelf smartphones to allow users to access and manage access to resources [4]. To be accessible by Grey, a physical resource, such as an office door, needs to be outfitted with a Bluetooth-enabled embedded computer and an electric strike. Unlike a system where all access-control policy is managed from a centralized location, Grey enables each user to delegate the authority they have to others, at their discretion. In this way, access-control policy is managed in a distributed fashion.

More specifically, each credential—which is a statement of authority (e.g., a delegation)—is expressed as a digitally signed certificate. Certificates are created and managed on phones; they are not stored in any central location. To access a resource, a Grey user instructs her phone to send to the resource, over Bluetooth, a set of credentials and a proof showing how that set of credentials fulfills the resource’s access-control policy. This access-control policy contains a nonce to protect against replay attacks, and the resource informs the user of its policy at the beginning of each interaction. The credentials, the access-control policy, and the proof are specified in a formal logic. Additionally, the phone requires that a PIN

²The wireless phone market is projected to reach 3 billion connections by the end of 2007 [24].

be entered before the phone can be used to exercise the user’s authority, e.g., before accessing a resource or issuing a delegation. For convenience, the PIN needs to be entered periodically, rather than for every access. If the phone were lost or stolen it would only be usable until the PIN timed out. Although these and other details are important to ensure the soundness of the system, they do not directly affect the user experience so we will not discuss them further here.

In a typical access scenario (which is usually any but the first attempt to access a resource), a Grey user causes a resource to open by selecting the Grey application on her cell phone, and scrolling to and selecting the list item representing the resource she wants to access (as shown in Figure 2). One could imagine different ways of initiating each access, e.g., by pointing the phone at the door to be opened. Some of these alternate methods aren’t feasible given the technological limitations of off-the-shelf smartphones. Also, requiring the user to select the resource to be accessed from a list has some benefits. For one, it is difficult to inadvertently cause a resource to be accessed; this is particularly relevant when, for example, the user could access any of the offices along a long corridor, but wishes to access only a specific one. Another advantage is that by clicking on a menu item it is possible to initiate an access when the user is not in close proximity to the resource. For example, some users initiate an office-door access as they leave the garage or enter the building, after which they put away their phones; when they reach Bluetooth range of the door (around 30 feet), the door simply unlocks. Finally, users who often access several resources in sequence can choose to have the entire sequence of accesses started by a single click, after which the phone will automatically access each resource in the sequence as the user approaches it. A common sequence is to unlock a perimeter door, then an office door, and then log into a computer.

A key feature made us choose Grey over other systems is its ability to support dynamic, end-user-based policy creation. Unlike smartcards, swipe cards, and RFID tags, which require that access-control policy is centrally specified by an administrator, Grey allows end-users to create and modify policy as they see fit, as long as the new policy is consistent with the rights that the users have been granted. Specifically, Grey users can create and modify policy using their smartphones either proactively (Figure 1), by using the Grey address book or a wizard interface to assign rights to users with whom they have previously interacted, or reactively, in response to another user’s access attempt that cannot succeed until that user is granted more authority. The typical ways of delegating authority are to (1) allow one-time access, (2) delegate only the authority to access a single resource, and (3) delegate all the authority possessed by a user. Grey supports the use of groups and roles. For example, Alice can group her students into a group called “Alice’s students” to allow her to more easily give all of them access to her



Figure 2: A Nokia N70 displaying the Grey resource list.

lab. The type of delegation that conveys all authority is normally used only with groups and roles (e.g., Alice extends to each of her students all the rights encapsulated by “Alice’s students,” but would delegate her own rights more restrictively). Once created, delegations are transferred between users using the cell phone network.

To better understand a typical Grey reactive delegation, imagine that Alice needs to get into Bob’s office. Alice selects the Grey application on her cell phone and selects Bob’s office, as shown in Figure 2. Her phone then contacts an embedded computer governing access to Bob’s door.³ Since Alice’s phone does not yet have sufficient credentials to access the door, her phone prompts her to ask someone for a delegation. Suppose Alice selects Bob and sends him a request for access to his office. Upon receiving this request on his phone, Bob can choose to give Alice either a short-lived credential valid for one access, or a more permanent delegation. In this way Bob can construct policies that allow him to grant access easily to multiple resources at once and to authorize an entire group of people for additional resources. Once Bob creates a delegation or denies the request, a message is returned to Alice’s phone which either carries credentials enabling her to unlock the door or notifies Alice that Bob has denied her request.

It is not difficult to imagine that this kind of spur-of-the-moment policy-creation ability allows users a great deal more flexibility in forming their policies than is provided by most access-control systems. This flexibility may not always be appropriate (e.g., the policy at a military installation may purposefully not permit modification except by administrators), but when appropriate, such as in university environments where end-users often have rapidly changing information access needs, it can be very convenient [13]. Although the technologies that underlie Grey are not the only way to achieve this flexibility, any end-user-based system that allows dynamic policy creation will need devices with displays, keypads, and the ability to communicate with infrastructure, and is thus likely to encounter user-interface and usability issues similar to the ones we discuss in this paper.

2.2 Related Work

The area of usable security is still a relatively new field of research. While security administrators may be well versed in security, end users do not have the technical experience necessary to make complicated security decisions [1, 12, 26]. When forced to make such decisions or to work under cumbersome security policies put in place by administrators, users tend to make poor decisions or even find ways to circumvent the system [1]. Thus end users can easily become the weakest link in a secure system.

Only a few published studies have examined the usability of authentication tokens such as smart cards or key fobs. One study found that most authentication tokens are not very usable, and those that are more usable tend to be less secure [7]. Another study found that seemingly simple authentication tokens can be difficult to use in practice. For example, smart card users often required several attempts to figure out which way to insert the card into the reader [23]. These studies look at the usability of authentication tokens themselves but do not consider how these tokens or the rights they convey are created or distributed between users.

There has been some work on user-interface design related to distributed access control for file systems. Cao showed that standard access-control list (ACL) interfaces had a high failure rate, despite users expressing confidence that they had manipulated the ACLs accurately [9]. Other studies showed that low levels of feedback in many access-control systems make it difficult for users to

³In our present deployment, this is a computer embedded in the wall that controls the electric strike on Bob’s door.



Figure 3: Floor plan of the 2nd floor of the office building. Grey-enabled doors used in this study are circled. Arrow points to kitchenette door.

understand what is wrong and what needs to be changed [19, 21]. Users also have difficulty understanding how different policies interact; for example, when a group is granted access to a resource but an individual who is a member of that group is denied access [21]. These studies look at how users build and manipulate access-control policies. However, they don’t take into account the restrictions imposed by small screens or other factors unique to mobile devices.

Several studies have investigated the usability of PGP and public key cryptography. Users have difficulty understanding the concept of certificates and public/private key pairs and as a result have difficulty doing such simple tasks as signing and encrypting email, let alone verifying the authenticity of a message [14, 26]. However, even when users understand how to use an encryption tool, they may fail to use it due to social concerns such as the fear of being perceived as paranoid [15]. The difficulty users have understanding public key cryptography concepts is relevant to our study of Grey; however, these studies focused primarily on the use of cryptography to secure and authenticate email.

There are several proposed distributed systems that use portable devices to control access to physical spaces [6, 27]. However, as far as we know this is the first published usability study of such a system.

3. METHODOLOGY AND DATA

To determine the reasons for users’ acceptance or rejection of Grey, we conducted a user study in which we observed a group of users as they transitioned from a security system based on keys to one based on Grey. We conducted interviews, logged Grey usage, and videotaped certain activities.

3.1 Environment

The study was conducted in an office building at Carnegie Mellon University. Over three dozen doors were connected to computers embedded in nearby walls. Each computer was able to communicate with the phones and to lock or unlock the door. The Grey phones and doors were set up to log all Grey-related activity.

The second floor of the building includes a large workspace that is locked after 6 P.M. and on weekends. The workspace has five perimeter doors, all of which were Grey-enabled. Inside the perimeter is a common area with a large number of cubicles for students and staff. On the edge of the common area are offices, labs, and storage

rooms used primarily by faculty and staff. We Grey-enabled 11 of the offices, two storage closets, one conference room, and one lab. In addition, we Grey-enabled a large machine room located on the first floor of the building. Several other doors in the building were Grey-enabled, but were not used in the course of this study. One of the perimeter doors can be unlocked only using Grey, while all the other Grey-enabled doors can also be unlocked using traditional keys.

3.2 Users

In January 2006 we began distributing Nokia N70 smartphones (shown in Figure 2) with Grey software installed. The users who received Grey phones were selected from faculty, staff, and students who either had a desk in the office building or had a regular need to access Grey-protected areas. The number of potential participants was limited in part by the willingness of users to switch to a phone that would support Grey. Similarly, many potential participants were turned away because they worked in buildings in which it would have been difficult to outfit offices with Grey. We tried to select users who were in working groups with other Grey users to maximize the usefulness of Grey.

We initially handed out Grey phones to only a few users. As the system became more stable and usable we increased the number of users incrementally. By the end of June 2006 we had 19 Grey users participating in our study. At the time of this writing all study participants had been using Grey for at least three months. One additional user participated briefly before dropping out of the study. In addition, Grey is used actively by the four authors of this paper and five other Grey project members.

The 19 Grey users participating in the study include 6 computer science and engineering faculty members, 9 computer science and engineering graduate students, 3 technical staff members, and 1 administrative staff member. 16 are male and 3 are female. To preserve privacy we refer to Grey users by fictitious names.

3.3 Procedure

Before giving a Grey phone to users we conducted an initial interview that explored their current security practices and how they managed their physical security in the office setting. The purpose of this interview was to understand the users' current work practices and concerns as they related to office security. If a user did not have an office we asked about other locations, such as their home. The primary focus of this study was to understand how users managed their own security.

Each user was then given a Grey phone and basic instruction on how to use it. We showed them how to open a door and request access from another person. We also informed them that if they became too frustrated at any time or if Grey failed to work it was perfectly acceptable to unlock a Grey-enabled door with a key.

After one month each user was interviewed again with the goal of understanding their initial use of Grey. This interview explored the user's use or lack of use of Grey's features as well as problems they encountered. We also asked how and why each user made use of Grey's delegation capabilities.

For the remainder of the study we interviewed each user every four to eight weeks, depending on user availability and activity. Interesting events, such as delegations, happen rarely so we used log data to determine when to schedule interviews. During these interviews we asked questions to determine how each user's interactions with and attitudes about Grey were changing over time. We also asked them about any changes they made to their access-control policies.




Figure 4: Diagram of the kitchen area where the timing study was performed. The camera was set up so that it captured the door and as much of the rest of the kitchen as possible.

All interviews were conducted at the participant's desk or in a nearby conference room. We wanted to make the interviews as contextual as possible and encouraged participants to show us artifacts such as their key chains or the contents of filing cabinets. Additionally, all interviews were audio recorded for later reference.

3.4 Videotaping

As a sub-study within our larger usability study of Grey, for two weeks we videotaped users unlocking a single highly trafficked perimeter door to better understand the differences between the way Grey and key users open a door. This door (shown in Figure 4) was located near a kitchenette and restrooms that were outside the locked workspace. People regularly used this door to return to the workspace after visiting the kitchenette, restroom, or other areas of the building.

The camera was set up 20 feet from the door in order to give it the maximum viewing range possible. People entering the kitchen from the hallway entered the camera's range approximately 14 feet from the door. People who entered through the Grey-enabled door and made use of the kitchenette remained within the camera's range until they left.

Key users were recruited by sending out a general email notifying them about the study and by asking them to participate as they passed through the door. During the study we turned off the camera when those who declined to participate accessed the door.

We videotaped door accesses for two hours every evening for two weeks. A total of 18 users were taped. Five Grey users accessed the door a total of 17 times and 17 key users accessed the door a total of 53 times. Some Grey users accessed the door with both keys and phones.

3.5 Videotaping Coding

In order to make our observations of key and Grey accesses comparable we picked several events that were logically similar in both processes. The events and the average times between them are shown in Figure 5. *Getting token* is when the user reaches for their phone or keys. *Stop at door* is when the participant stops in front of the door or when he approaches within arms reach of the door and significantly slows his speed. *Door opened* is when the door is unlatched and pulled forward and *door closed* is when the door closes again and the latch clicks into place.

The events don't necessarily occur in the order listed above. For example, a person may stop in front of the door before reaching for his keys. In this case we included the time between reaching for the key and getting it as between the get-token and stop-at-door events. However, we only recorded six cases of this and it had minimal impact on our results.

3.6 Data

The study lasted for one year. During that time our logs recorded 19,500 Grey access attempts and 236,900 total door accesses including both ingress and egress. While monitoring the kitchenette door we videotaped 70 accesses, 17 of which involved Grey. Finally, we audio recorded approximately 30 hours of interview data.

Individual users made frequent use of Grey. When users made use of the system⁴ they averaged 12 access attempts per week. This ranged from heavy users who averaged 35 accesses a week to light users who only used the system a few times a month. Since the perimeter doors were unlocked during business hours users with offices used Grey significantly more often. Five of the users accessed their offices almost exclusively with Grey. Three users gave away their keys because they had no longer needed them. Most participants' use of Grey remained relatively stable over the course of the study.

Each phone maintains a list of the resources (e.g., office doors) that it has been used to access, so that those same resources can be conveniently accessed in the future with a single click. The average number of resources on a given user's list was 7.4, with a maximum of 15 and a minimum of 2.

The Grey application on each phone also contains a Grey address book which contains the names, numbers, and public keys of other Grey users. Address-book entries represent users to whom access can be delegated and who can be asked to facilitate an access. An address book can be populated with entries via, e.g., a phone-to-phone business card exchange. Users in our study had an average of 5.7 other users listed in their address book with a minimum of 3 and a maximum of 11.

Of the 19 users who began the study, 18 are still actively using the system despite being given the option to stop at any time. The remaining person has kept the system on his door but doesn't use the system for opening doors. A user whose office is located in another building even offered to pay to have the system installed on his office door. Several other users expressed interest in having such a system installed in their home.

4. LESSONS LEARNED

After collecting and analyzing the results from the interviews, logs and videotapes we found several different reasons why users rejected or accepted the new technology. Each reason is an instance of a more general principle that manifested itself in our system in a specific and sometimes unexpected way. We detail these principles and how they applied to Grey in this section, and from each we attempt to draw lessons to aid in the design of other access-control technologies as well as smartphone and mobile-device applications in other domains.

Principle 1: *Perceived speed and convenience are critical to user satisfaction and acceptance.*

The designers of access-control systems typically focus on the security properties of such systems and their ability to support a va-

⁴We are excluding weeks when the user made no use of the system. Events such as conferences, vacation and holidays caused our users to be absent at unpredictable times and therefore unable to use the system.

riety of access policies. However, we observed that end users tend to be most concerned about how convenient they are to use. There are many examples of end users of widely used access-control technologies readily sacrificing security for convenience. For example, it is well known that users often write their passwords on post-it notes and stick them to their computer monitors. Other users are more inventive: a good example is the user who pointed a webcam at his fob and published the image online so he would not have to carry the fob around.⁵

The Grey users in our study never raised any concerns about the ability of Grey to provide adequate security. However, we received many complaints about the speed and convenience of accessing resources with Grey, and we observed users sacrificing security for convenience. This was an especially interesting observation given that many of our users do research in the computer security area.

The following anecdote illustrates how some Grey users sacrificed security for convenience. One evening we observed a Grey user taking a magazine off a nearby magazine rack and placing it in the doorway of a perimeter door to prevent it from locking when he left the workspace briefly. When asked about the incident he pointed out that the clothing he was wearing at the time had no pockets making it inconvenient to carry any objects. He also pointed out the relative insecurity of the perimeter doors, noting that pizza delivery guys easily gain access to the locked workspace on a regular basis. Given this reasoning, we expect this Grey user probably behaved in a similar way when he was using keys. Both Grey and keys require that users carry a physical token that can be inconvenient to carry, and neither approach addresses the problem of users who don't respect the security policy for a shared resource (in this case the policy that the perimeter doors are supposed to be kept locked after 6 P.M.).

We began receiving complaints about the speed of Grey shortly after distributing the first set of Grey phones to users. Five out of eight initial users told us that they thought Grey was slower than their keys when we interviewed them a month after they received their phones.

We analyzed 335 door accesses in our log files and found that with Grey it took an average of 6.6 seconds (standard deviation of 1.7) from the time the phone first sent a request to the moment the door unlocked. This latency seemed acceptable since initial rough measurements indicated that keys were approximately as slow to use, and we anticipated that with Grey users would attempt to unlock the door while they were walking toward it. In order to understand why some users were dissatisfied with the speed of unlocking doors with Grey, we wanted to observe not only how long an entire access took, but also how long within that time it took for our users to interact with the system, as the effective speed of an access-control system is highly dependent on how users manipulate their access-control tokens [23]. Therefore, we used a video camera to record both key and Grey users accessing a highly trafficked door, as discussed in Sections 3.4 and 3.5.

The results of our observations are summarized in Figure 5. Briefly, opening a door (measured from the time a user reached for his keys or phone to the closing of the door after the user walked through) took roughly the same amount of time using Grey or keys, and parts of the process actually took less time with Grey. The time difference between keys and Grey for each of the three steps, selecting access token, unlocking door, and going through door, is statistically significant. How, then, to explain our users' impression that Grey was slow? Our analysis of the videotaped door accesses revealed a difference in how time was spent between key and

⁵<http://fob.webhop.net/>




Figure 5: Average door access times for key and Grey users who fetched their key (18 accesses) or phone (7 accesses) in sight of the camera. Most notable is that Grey users spend more time waiting in front of the door but less time moving through it.

Grey users. Key users were always active while accessing the door; they were either finding, inserting, or removing their keys and never spent any time just standing and waiting. In contrast, Grey users only had to push a button and wait so they spent more of their time idle. In fact, Grey users spent an average of 2.9 seconds idly waiting outside the door, a time well above the point at which users begin to notice a delay. Nielsen explains that there are three time limits for human perception of system response time. A delay of 0.1 seconds or less will not be noticed; if the delay is under a second the users will notice but will not break their flow of thoughts; and, finally, any delay longer than 10 seconds risks losing the user's attention. [22].

Additionally, much of the time savings experienced by Grey users happened after the door was already opened. Grey users spent an average of 3.8 seconds going through the door compared with the 5.7 seconds experienced by key users (this is a statistically significant difference). Since the users' goal is unlocking the door, we hypothesize that users using keys do not perceive this extra time since their task, unlocking the door, has already been completed.

Human reaction to system response time is a well studied topic in HCI [17, 18, 22, 25]. Human perception of time is very different from clock time [16]. Users perceive "busytime when they are engaged in a task as taking less time than waiting for the system to do something. Users also do not notice the passage of time when they are engaged in high-level cognitive functions. For example, Tognazzini described user testing of the Apple Human Interface in which most users reported that keyboard shortcuts were faster than using a mouse, when in fact the opposite was true. Tognazzini attributed this to the fact that remembering the abstract symbols associated with a keyboard shortcut is a high-level cognitive function, while moving a mouse and selecting a menu item is not. Thus users were unaware of the time they spent remembering the shortcuts, but were keenly aware of the time they spent moving the mouse [25].

Another factor may be the social awkwardness associated with waiting in front of a locked door. Logan tells us how opening a door with Grey can be frustrating. Since some doors in the building are made of glass it is possible for those inside to watch someone outside trying to get in. Logan points out how this can be embarrassing when he has to wait for the door to unlock.

I yank on the door and am very surprised [that it is locked] for a couple seconds and then I find myself standing outside and everybody inside is looking at me

standing outside while I am trying to futz with my phone and open the stupid door.

We observed that many key users have gotten into the habit of pre-fetching their keys before they reach the door. Out of 30 videotaped accesses where the key user appeared from somewhere outside the camera's range (i.e., they did not use kitchen facilities or stop to talk to someone within the camera's range before unlocking the door) 26 showed the key user entering the camera range with their keys already in hand. In fact, many users seemed to have optimized their key-rings to accommodate finding the correct key quickly and easily while walking. Out of 20 participants interviewed 9 organized their key rings so they could quickly locate keys without needing to look at them.

Thomas explained how he was able to find his keys in his pocket with little to no effort.

There is no attention paid to keys. I mean, at this point I can generally, especially with my house key and my car key, there is like a better than 70% chance that if I want one of those keys I just like dig into my pocket and grab it out and I will actually have that key in my hand, um, just from feel of it.

With so many users pre-fetching keys it is our hope that with practice Grey users would learn to do something similar. And, in fact, in later interviews users began discussing how they had learned the exact place to contact the door so that it opened with a "satisfying click" as soon as they arrived. (A Grey phone can begin its dialog with a door as soon as it is within Bluetooth communication range.) Anthony explains how Grey is faster using this approach but only if he remembers to select the door in time so that all the access time occurs while he is walking.

I could push the button while I was walking down the hall so it was open by the time I got here. Um, so as long as I remembered to get the phone out and push the button it was faster.

After receiving a number of complaints from users about the speed of Grey, we made some changes in the Grey software to improve performance. We were able to reduce the amount of time it typically takes the door to unlock by 2 seconds. Since Grey users waited in front of the door for an average of 2.9 seconds, this was a significant improvement in the user experience.

We interviewed Anthony shortly after updating Grey to the new, faster version.

I don't know what the new timing is now but it is fast, it's faster than keys now.

Users want to access doors very quickly without having to think about what they are doing. Key users have perfected their technique over time by re-organizing their key ring and training themselves to use keys quickly with little concentration. Grey users can similarly benefit from optimizing their Grey usage by activating Grey before arriving at a door. Activating Grey early means less time spent standing around.

Our field study confirmed that users are often more concerned about speed and convenience than they are about security, even when using a security technology. This is likely related to the fact that security is a secondary task for most people [1]. It is critical that security-related technologies that require end-user interaction be convenient to use so that legitimate users do not try to circumvent them.

A more surprising observation was that Grey users perceived time spent waiting for Grey to open a door as significantly longer than time spent manipulating keys—even when the actual time usually differed by less than a second. One reason for this is that people don't notice how long it takes them to manipulate their keys because they are actively engaged in the process. Another reason is that people feel uncomfortable being observed standing in front of a door not doing anything. We expect that passive waiting time and the social stigmas associated with it may be a problem in other mobile applications as well. When application developers are unable to eliminate waiting time they should consider where people are likely to be and what they are likely to be doing during waiting periods, and attempt to move waiting time to parts of the interaction where it will be least conspicuous.

Principle 2: *A single failure can strongly discourage adoption.*

A single failure can cause the new adopter to lose faith in a new technology and revert to a more familiar approach [2, 11]. This is especially true in systems such as Grey where the associated cost of using the other system is low and the perceived potential cost of using the new system is high [5]. This problem is especially acute with Grey, because failures are expensive—the likely result of a failure is that a user will be locked out from his office or the floor.

While Grey is relatively reliable, it is not as reliable as a production system. Failures can occur for a number of reasons: delays or data loss in the cell phone network, firmware or operating-system errors on the phones themselves, bugs in the Grey software, door hardware or software failures, and misconfiguration or user error.

Anders describes how getting locked out one night due to a failure caused him to drop from an average of 28 Grey accesses a month to seven.

I've been using the phone regularly up to one point But then there is one time it breaks then, you know, it shatters my confidence. So from then on I stopped using the phone. Once it has proven itself not reliable then ... there is no added advantage for me to use it.

The cost of a failure is different depending on the circumstances under which it occurs. When Anders became locked out it was a devastating failure because it was late at night and there were very few people around. Zack had a similar experience that prevented him from opening his office; however, since it happened during business hours, he simply borrowed a key and let himself in. He wasn't overly bothered by the experience; in fact, his use of Grey steadily increased and he continued to leave his keys in his office.

A user can also lose faith in the system if he perceives something as “not working,” even when there is no failure. Notably, requesting and receiving authorizations via delegation in Grey was sometimes so slow that users perceived them to be broken. Donald tells us about how he asked another user for a delegation using Grey but was forced to cancel the request when it took too long.

When everything is working it's ... it's OK but it's like the failures that it has, um, that like especially that there is no feedback.... I wasted a lot of time just waiting for something to happen on there and eventually I just like put the phone in my pocket and I went over to [another building] to do something else and when I came back and it was still like “waiting for answer” or something like that.

When a user requests an authorization from another user the message is sent over the wireless service provider's SMS and GPRS networks; this makes it possible for two users to communicate regardless of where in the world they are or how distant from one another. Unfortunately, the SMS networks are occasionally very slow and it can take a while for messages to go through. Furthermore, once a message is received, there is no way to be certain that the recipient will notice it (e.g., he may not be in possession of his phone, or his phone could be in silent mode).

This happened when Riley needed to get something out of May's office one day when she was working from home. Riley called May and explained the situation and after hanging up sent a Grey request. Nothing happened for several minutes so he tried sending another request and again nothing happened. Finally, he called May again and had her create an authorization and proactively send it. The authorization arrived after several minutes. Talking to May, we learned that she eventually received one of the requests, over 15 minutes after it was sent. A few weeks later when May stopped by Riley's office to ask him to delegate access to her for a room she needed to access briefly, Riley instead handed her his Grey phone and told her she could use it to let herself in. When asked about the incident, Riley said he was concerned that the authorization request would likely fail as it had in the past and it would be much faster to just lend out his phone.

In general, designers of access-control systems and other security products focus primarily on keeping people out, and only secondarily on making sure they can get in. For some environments, this is a mistake, as the consequence of a failure in allowing a person to access a resource can be more dire than the consequence of erroneously allowing access. Another lesson we draw from this experience is that it should always be possible to distinguish the correct (but perhaps undesired) behavior of a system from a failure, and that, if possible, the user should be informed of the nature of a failure so that he can judge the likelihood of its reappearance. Along these lines, we augmented our system to inform the requester of a delegation whether his request was received (as opposed to still in transit) as well as whether it was acted on (seen and acknowledged by the recipient). We also undertook efforts to recognize and inform the user of errors that likely resulted from misconfiguration. We have yet to determine to what degree these measures increase users' confidence in the system.

Technology failures are always discouraging, but computer users have come to expect some failures and have become somewhat tolerant of them. For example, personal computer users have become used to the fact that applications “hang” and that they need to frequently reboot their computers. Our field study suggests that users may be less tolerant of failures in mobile devices or access-control technology, especially when there is not an easy way to recover. So-

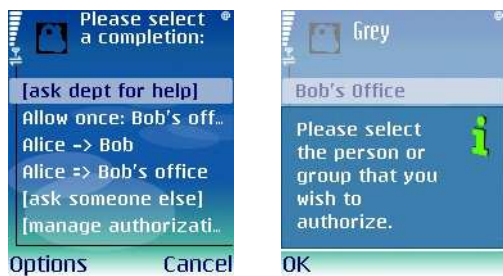


Figure 6: Left: Old reactive delegation interface. Right: New proactive delegation interface.

me Grey failures required users to get an administrator to “reboot their door.” Users had a low tolerance for this when it occurred during working hours, but users who experienced this problem late at night found it completely unacceptable. Grey could be improved by both making it more stable and giving the user more feedback on the state of the system.

Principle 3: Users won't use features they don't understand.

Users are reluctant to use options provided by the interface when they don't properly understand the consequences. They are most likely to pick the option they understand the best, even when they know it is not the option they want. This is exacerbated in situations when users do not know whether they can backtrack after making a selection [5]. In our study we witnessed users passing up more effective methods of delegation for less effective but simpler-to-understand methods.

When reactively creating a delegation in response to another user's access attempt, a Grey user needs to choose among a set of possible delegations computed by the phone. For example, if Alice is asking Bob for help, Bob could delegate directly to Alice, or he could delegate to Charlie if he knows that Charlie has already delegated to Alice. The delegations could also convey different levels of authority, involve indirection through groups, etc. In a system populated with credentials, any of typically at least a handful of different delegations can satisfy an incoming request.

The initial interface design, shown in Figure 6, attempted to present all the possible delegations to the user as a list. On a full-size computer screen this could probably be managed in an understandable way with any of a number of interface designs. The Nokia N70, however, has a screen resolution of 176x208 pixels, which allows only about 20 characters to be displayed on a single line. With so little screen real estate it becomes very difficult to display all the relevant information to the user at once. Hence, we introduced abbreviations to describe the different kinds of delegation, and decided to forego including any instructions on the interface itself. Even so, delegations were often too long to fit on one line and scrolled off the screen (though the full line could be seen with two additional button clicks). We believed that users would be willing to learn the abbreviations (there were only two) and put up with the brevity of the representation; in exchange, they would have the convenience of being able to answer a help request simply by scrolling through and clicking on an item on a list.

It quickly became clear that this interface was not meeting users' needs. Of the five users who created delegations in the first month of use, none actually knew what all the different options meant. For example, after observing Riley responding to a request from Donald we asked why he had selected “Allow Once” as opposed to

giving a longer delegation. He replied that the “Allow Once” option was the only one he understood.

An obvious solution was to re-implement the interface as a wizard. A wizard is a user interface that constrains the user to doing a task one step at a time in a specific sequence, often stepping through several screens to complete the task. Though wizards are very useful in many situations, one of their biggest pitfalls is that they can unnecessarily make a short task much longer by forcing the user to go through multiple screens instead of just one. This, in fact, was the main reason why we didn't use a wizard in the first place.

However, after we implemented a wizard interface for proactive delegations and got a positive reaction from users (the interface problem described here relates to reactive, rather than proactive, delegation), we decided to do the same for reactive delegation. Using a wizard, the user could specify each different part of the delegation (e.g., what kind of delegation, to whom) on a separate screen, with a screen of instructions preceding each input screen.

We built a small paper prototype of the new reactive delegation interface and asked several users to respond to a request from Bob to get into their office. If they elected to “Allow Once” we asked them to assume they had a good reason to create a longer delegation. We got a very positive result: all eight users were able to successfully create their intended policies and were able to understand all the options. However, five of the eight users asked still initially selected “Allow Once.” Anthony explained that even though he knew Bob he didn't want to give him access without talking to him first.

I want to have a conversation with [Bob] before I give [access to my office] to him for all time.

In summary, we learned that even technically savvy and motivated users as a rule were not willing to put much effort into learning how to use a concise but not immediately clear interface. Providing cues to users about how to use features is particularly problematic on hand-held devices with small screens. Grey users told us they were ignoring most menu options because they did not understand them. Since the small screen precluded adding clarifying words to the existing interface, we switched part of the system to a wizard interface that broke a task up into several small steps. This may be a good solution for other mobile applications as well.

Principle 4: Systems that benefit from the network effect are often untenable for small user populations.

A system benefits from the *network effect* if the addition of a new user causes existing users to get extra benefit from the system. Such a system becomes truly useful when it accumulates a critical mass of users; conversely, the system can be of little utility to its users until critical mass is achieved [20].

One of the most potentially useful things about Grey is the ability of users to spontaneously give out delegations to each other. One of the reasons Grey was designed to work on off-the-shelf smartphones was to increase the potential user base. Unfortunately, our software does not yet run on very many kinds of phones, and thus currently only the 19 users in our study and the 10 Grey project members are able to delegate to each other.

Some people with offices or cubicles in our building could have benefited from using Grey phones but were unable to participate in our study because they subscribed to a cell phone service that was incompatible with the Nokia N70 phone. Because only a fraction of the people in our building have Grey phones, some of our study participants found they had little need to delegate. Those participants who did not have offices often had no resources worth delegating

since they typically had access only to common areas readily accessible by all the participants.

Moreover, users recognized that the small user base meant that the utility they would derive from Grey would be limited, and this perception discouraged them from pre-configuring their Grey software to make delegations easier to issue. Ironically, this meant that once there was an opportunity or need to delegate, the cost of delegating was higher.

Anthony explained why he had not added anyone to his Grey address book, the first step in proactively delegating access.

I haven't because I am only working with one person that has a Grey phone right now and he sits in a cube. ... I didn't see any reason to add him. ... Since he doesn't have an office what would I gain by adding him to my address book?

During the study we noticed many occasions when Grey could have been very useful, but unfortunately the potential recipient of a delegation did not have a Grey phone and would only have needed one for that one occasion. For example, Grey users sometimes had visitors come to their offices after the perimeter doors had been locked. Anders mentioned that he conducted user studies in the second floor lab on weekends and had to wait by a perimeter door to let participants in.

Unfortunately, this bootstrapping problem cannot be solved easily, though steps can be taken to prevent users from becoming discouraged by the apparent lack of situations in which the new technology can be useful. When there are start-up costs that are typically amortized over many uses of a technology—as is the case with filling a Grey address book with potential recipients of delegations—these costs could instead be paid up front or amortized by other benefits (e.g., a prize for the most active user). That way when an actual need to use the system arises, all the start-up costs have already been paid, minimizing the overhead to using the system. The network effect may also be dealt with by making non-Grey users more active in the system. This could be done by allowing Grey users to unlock a door remotely, much like a buzzer system in an apartment building, so that non-Grey users can benefit from the system.

There are many technologies that are of limited use in a vacuum. In order to do interesting things with them you need to know other people who own interoperable devices. To get the most use out of a Grey-like system, all people who interact with Grey-enabled doors should have Grey-enabled phones. Unfortunately, limitations in currently-available smart phones and budget restrictions made it difficult for us to make Grey available to as many people as we would have liked for our field study. Developers of mobile applications can reduce this problem by developing code that depends as little as possible on specific hardware platforms. Hopefully, as smartphone technology improves this will become less of a problem. In the meantime, field studies might include incentives to help bootstrap use of the system.

Principle 5: *Low overhead for creating and changing policies encourages policy change.*

One of the main goals in designing Grey was to enable users to create access policies that are more flexible, convenient, and secure than the policies designed to be enforced by keys. We explore this topic in depth in a separate work [3] and discuss it here only briefly.

The pre-Grey access policy in our building illustrates the very coarse granularity of most key-based policies. For example, a role-based key policy resulted in students, staff and faculty being given different keys. Student keys opened a minimum number of

doors, because students were judged less responsible than faculty and staff. Ethan, a student, told us how he needed access to one supply closet every couple of weeks but since he wasn't given a key he had to go find someone to let him in every time. He could not obtain a key because the key to the closet also opened other areas to which he was not permitted.

In addition to not always allowing the enforcement of the desired policy, obtaining new keys is typically a time-consuming process, making it inconvenient to use keys when on-the-fly delegation is needed. Sara told us how she kept a spare set of keys in her office to lend to temporary employees while they waited to get their official keys. In some cases, it could take more than a month to obtain keys for a new temporary employee. Amy had a similar story: she too kept extra keys to lend while official ones were being requested. She once accidentally gave a new employee the key that opened her office as well as the outer doors. Though she was able to recover the key, the experience reminded her how careful she had to be when lending keys; the one she had mistakenly lent out could easily have been copied.

The inability of keys to express precise or ad-hoc delegations gives rise to the use of *hidden keys*, keys hidden in public locations where they are available for group or emergency use. There were three sets of hidden keys maintained by users in our study. In each case the hidden keys were used as a way to allow all members of a group occasional access to an area without giving each member of the group an individual key. The use of these keys was unregulated and it was often unclear who knew about the hidden key or who was using it.

Thomas tells us that the shared key works on the honor system and that occasionally keys would go missing:

People just come and take [a shared key] and, um, people are very good about bringing them back. ... Once or twice people have, including me, has accidentally taken [a shared key] home with them or something like that. And I send out a mail like, um, could someone please bring back this key.

After its introduction, delegations made through Grey started taking the place of hidden keys. In some cases, people who had access to doors via the hidden keys were not issued a corresponding delegation in Grey. In most cases, this was because the users had access via the hidden key as a side-effect of the clumsiness of keys, rather than as part of a desired policy.

For example, with keys and hidden keys Anthony had little to no control over who could get into his lab. As a consequence, Mark, who knew where a set of hidden keys was located, was given inadvertent access to the lab. However, Anthony saw no reason for Mark to access the lab, so he didn't issue Mark a corresponding delegation via Grey. Thus, the policy enacted by Grey was more secure (from the standpoint of Anthony) than the policy implemented by keys.

With Grey, users also began to delegate access more casually. Of our 19 users, 11 received delegations to resources that they previously could not access. For example, Ethan had an occasional, but not very pressing, need to access Keith's office. Ethan hadn't previously had a key to the office because Keith judged that it wasn't worth the effort to procure him a key. With Grey, delegating was sufficiently easy that Keith immediately delegated access to Ethan. Keith describes his reasoning:

[Delegation with Grey] was easy. Getting a key for him would not have been easy. I don't know, it just sorta came up, now that [I] am using Grey [I] can do this kinda thing.

By allowing easy-to-implement, fine-grained control, Grey empowers users to make new policies that better fit their needs. In a separate paper we provide an in-depth analysis of how policies implemented using Grey compare with users' ideal policies and existing policies implemented using keys. We interviewed users to determine their ideal access control policies for Grey-enabled resources they control, as well as the the actual access control policies they had implemented using existing keys. We used Grey log file data to determine what policies users implemented with Grey. We found that Grey policies correctly matched the user's ideal policy for 95% of the access-control rules. In all other cases the ideal policy specified capabilities not currently implemented in Grey. In contrast, existing key policies correctly implemented between 30% and 92% of the ideal access-control rules, depending on what assumptions were made about who had access to hidden keys [3].

When a new technology makes it easy to do something, it is likely that people will do it—but only if it was something that they were interested in doing in the first place. Prior to Grey, creating and changing access-control policies in our building was extremely difficult, requiring that new keys be made, old keys be returned, or locks be re-keyed. Even if the distribution and creation of keys was optimized it would still be inconvenient for users since at least one third party, the locksmith, would need to be involved. Grey users in our building were able to quickly and easily give short-term or long-term access to other Grey users without needing to involve anyone else. We observed Grey users creating new policies that would not have been worth the effort for them to implement without Grey. Our field study suggests that people have needs for access-control policies to physical spaces that are difficult to implement with keys, and that there is a need for a system like Grey that enables more flexible policies.

Principle 6: *Unanticipated uses can bolster acceptance.*

Unanticipated uses are a good sign for a system that potentially suffers from the network effect. Increased usefulness bolsters acceptance which in turn encourages more people to join the system. The unexpected uses our participants found for Grey show how such a system can have more value than as simply a replacement for physical keys. Future designers should keep these types of use cases in mind to ensure that their user interfaces are sufficiently flexible to allow unintended uses of new technology.

One unanticipated use of Grey was unlocking office doors from the inside. Eric commented that his favorite part of Grey was that he no longer needed to get up from his desk to open the door to let someone in. He simply unlocked the door with his phone and told them to come in. This was a very useful feature to him, because during meetings when his office was full it could be difficult to find a path to walk to the door. Additionally, he found getting out of his chair to open the door disruptive to his work.

Being able to unlock a door from a distance is useful in other situations as well. While doing the videotaping discussed earlier we watched a few Grey users participate in a group dinner in the kitchenette. For various reasons, different members of the group needed to go in and out through the locked door. Eric quickly realized the inefficiency of using keys and simply put his phone on the table. Every time a group member headed for the door he would hit a button and the door would unlock.

Other users discussed how enjoyable it was to surprise friends by unlocking the door from a distance. Logan pointed out how the phone was a “cool new toy” making it fun to play with. He also commented on the “satisfying” clicking noise the door made when it unlocked.

System designers often focus their design efforts on a “killer app,” or application that best demonstrates the advantage of their system. In doing this they often neglect, or even fail to think of, other modes of use for their system, which, even if trivial or tangential to the main intended use, nevertheless have the potential to strongly encourage users to adopt the new technology. We were surprised by some of the uses people made of the Grey system. Using a Grey phone as a remote control to unlock a door from a distance was one such unanticipated use. Now that we realize that people want to use Grey in this way, we will consider interface changes that will further encourage this use. Without a field study we would have been completely unaware of this use of our system.

5. DISCUSSION

Our observations of users in our study crystallize into several intertwined themes. The first theme is that users overwhelmingly treat their Grey-enabled phones as appliances. Users expected Grey to simply work—failure, latency, and inconvenience were not tolerated. We hypothesize that these expectations carry over from users' previous experiences with mobile phones and keys, which have simpler capabilities and fewer failure modes, but we were surprised that the additional functionality offered by Grey didn't engender greater tolerance for some tradeoffs. It is notable that users are typically willing to tolerate some level of undesired latency and failure (e.g., the need to reboot) when using their personal computers; this tolerance did not extend to what are, essentially, mobile personal computers.

The second theme has to do with user interfaces for taking advantage of the advanced functionality that Grey makes possible (e.g., creating policies with groups and roles). Although interviews with users indicated that many were interested in advanced features (e.g., most users used groups and roles when describing their ideal access policies), in practice almost all users simply wanted to achieve the desired effect (for a particular door to open) as quickly and with as little user-interface manipulation as possible even if this meant giving up on more complicated goals. Our initial interfaces, which were designed as a compromise between maximum efficiency and richness of features, satisfied almost nobody—most users wanted more streamlined interfaces, and a minority needed more powerful ones. The best course may be to design completely separate interfaces for the most basic and for advanced functionality, despite the claim of most users that they desire more than just basic features; we are currently pursuing this course with Grey.

The third theme, which emerged as being interwoven with the first two, is that the education, interests, and skills of a user population do not necessarily affect the kinds of user interfaces that the population finds useful. In particular, our users were on average extremely well-educated, tech-savvy gadget lovers, yet they showed little inclination or ability to learn to use any of our less than completely intuitive interfaces. Most understood the complexity hidden away behind the user interface, and appreciated the benefits that it could offer, but nevertheless had little patience for latency or inconvenience and therefore did not take advantage of Grey's features as much as we had anticipated they would.

Taken together these themes and the six principles discussed earlier demonstrate the importance of emphasizing ease of use and convenience in smartphone and mobile-device applications—even if they are targeted at “power users.” Unintuitive interfaces, system failures, and latency caused users in our study to avoid using some of the features provided by Grey, or to avoid using the system completely. However, we saw that when users found Grey features convenient to use, they took advantage of the flexibility it offers to

create access-control policies that more closely matched their needs than was feasible using keys.

While most of our participants were technically educated, we also observed users, some beyond the scope of this study, who were not. Most of our findings apply to the non-technical users as well as our technical ones. Indeed, we had expected to find our technical users would use more of Grey's advanced features and would not have the kinds of difficulties that we expected from non-technical users; instead, we found that our technical users generally behaved the way we expected less technical users to behave. Overall, we believe the lessons learned are generalizable to a wider population, though future studies with a more diverse set of users would help validate this.

Finally, many of the lessons we have learned from deploying Grey and the specific ways in which general principles manifested themselves in this deployment have already shown themselves useful in assisting with the design of other mobile-device-based systems, such as a people-finder application being developed at Carnegie Mellon [10].

6. ACKNOWLEDGMENTS

We thank all the students and staff that have worked to make Grey successful. We would also like to acknowledge the assistance of Rob Reeder in the writing of this paper.

This work was supported in part by National Science Foundation Cyber Trust Grants CNS-0433540 and CNS-0627513, and U.S. Army Research Office contract no. DAAD19-02-1-0389 ("Perpetually Available and Secure Information Systems") to Carnegie Mellon University's CyLab.

7. REFERENCES

- [1] A. Adams and M. A. Sasse. Users are not the enemy. *Commun. ACM*, 42(12):40–46, 1999.
- [2] A. Bandura. V.S. Ramachaudran (Ed.), *Encyclopedia of Human Behavior*, volume 4, chapter Self-Efficacy, pages 71–81. Academic Press, New York, 1994.
- [3] L. Bauer, L. Cranor, R. W. Reeder, M. K. Reiter, and K. Vaniea. Comparing access-control technologies: A study of keys and smartphones. Technical Report CMU-CYLAB-07-005, Carnegie Mellon University, 2007.
- [4] L. Bauer, S. Garriss, J. M. McCune, M. K. Reiter, J. Rouse, and P. Rutenbar. Device-enabled authorization in the Grey system. In *Proceedings of the 8th Information Security Conference*, pages 431–445, Sept. 2005.
- [5] L. R. Beach and T. R. Mitchell. A contingency model for the selection of decision strategies. *The Academy of Management Review*, 3:439–449, 1978.
- [6] A. Beaufour and P. Bonnet. Personal servers as digital keys. In *Proc. 2nd IEEE International Conference of Pervasive Computing and Communications*, Mar. 2004.
- [7] C. Braz and J. Robert. Security and usability: The case of the user authentication methods. In *IHM '06: Proceedings of the 18th International Conference on Association Francophone d'Interaction Homme-Machine*, pages 199–203, 2006.
- [8] Smart mobile device market growth remains steady at 55%. Canals Research Release 2006/071, July 2006. Available at <http://www.canalys.com/pr/2006/r2006071.htm> as of Sept. 23, 2006.
- [9] X. Cao and L. Iverson. Intentional access management: Making access control usable for end-users. In *Symposium On Usable Privacy and Security (SOUPS)*, 2006.
- [10] J. Cornwell, I. Fette, G. Hsieh, M. Prabaker, J. Rao, K. Tang, K. Vaniea, L. Bauer, L. Cranor, J. Hong, B. McLaren, M. Reiter, and N. Sadeh. User-controllable security and privacy for pervasive computing. In *Eighth IEEE Workshop on Mobile Computing Systems and Applications (HotMobile)*, Feb. 2007.
- [11] F. D. Davis. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3):319–340, Sep 1989.
- [12] P. Dourish, E. Grinter, J. D. de la Flor, and M. Joseph. Security in the wild: User strategies for managing security as an everyday, practical problem. *Personal Ubiquitous Computing*, 8(6):391–401, 2004.
- [13] D. F. Ferraiolo, D. M. Gilbert, and N. Lynch. An examination of federal and commercial access control policy needs. In *16th National Computer Security Conference*, pages 107–116, 1993.
- [14] S. L. Garfinkel and R. C. Miller. Johnny 2: A user test of key continuity management with S/MIME and Outlook Express. In *SOUPS '05: Proceedings of the 2005 Symposium on Usable privacy and security*, pages 13–24, 2005.
- [15] S. Gaw, E. W. Felten, and P. Fernandez-Kelly. Secrecy, flagging, and paranoia: Adoption criteria in encrypted email. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 591–600, 2006.
- [16] E. Geelhoed, P. Toft, S. Roberts, and P. Hyland. To influence time perception. In *CHI '95: Conference companion on Human factors in computing systems*, pages 272–273, 1995.
- [17] R. Geist, R. Allen, and R. Nowaczyk. Towards a model of user perception of computer systems response time. In *CHI '87: Proceedings of the SIGCHI/GI conference on Human factors in computing systems and graphics interface*, pages 249–253, 1987.
- [18] M. Hildebrandt, A. Dix, and H. A. Meyer. Time design. In *CHI '04 Extended abstracts on Human factors in computing systems*, pages 1737–1738, 2004.
- [19] A. Kapadia, G. Sampemane, and R. H. Campbell. Know why your access was denied: Regulating feedback for usable security. In *CCS '04: Proceedings of the 11th ACM Conference on Computer and Communications Security*, pages 52–61, 2004.
- [20] M. L. Katz and C. Shapiro. Systems competition and network effects. *Journal of Economic Perspectives*, 8(2):93–115, Spring 1994.
- [21] R. A. Maxion and R. W. Reeder. Improving user-interface dependability through mitigation of human error. *International Journal of Human-Computer Studies*, 63(1-2), 2005.
- [22] J. Nielsen. *Usability Engineering*, chapter 5. Morgan Kaufmann, 1994.
- [23] U. Piazzalunga, P. Salveneschi, and P. Coffetti. The usability of security devices. In L. F. Cranor and S. Garfinkel, editors, *Security and Usability: Designing Secure Systems that People Can Use*, pages 221–241. O'Reilly, 2005.
- [24] C. Taylor. Global mobile phone connections hit 2.5bn. The Register, Sept. 2006. Available at http://www.theregister.co.uk/2006/09/08/mobile_connections_soar/ as of Sept. 27, 2006.
- [25] B. Tognazzini. *Tog on Interface*, chapter 6. Addison-Wesley Professional, 1992.

- [26] A. Whitten and J. D. Tygar. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *Proceedings of the 8th USENIX Security Symposium*, 1999.
- [27] F. Zhu, M. W. Mutka, and L. M. Ni. The master key: A private authentication approach for pervasive computing environments. In *Fourth IEEE International Conference on Pervasive Computing and Communications (PerCom'06)*, pages 212–221, 2006.