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

Lujo Bauer, Lorrie Cranor, Michael K. Reiter, Kami Vaniea

October 18, 2006
CMU-CyLab-06-016

CyLab
Carnegie Mellon University
Pittsburgh, PA 15213

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

Lujo Bauer Lorrie Cranor Michael K. Reiter Kami Vaniea

Carnegie Mellon University
{lbauer, lorrie, reiter, kami}@cmu.edu

October 18, 2006

1 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 (i) aspects of the system that gave rise to failures, misunderstandings, misperceptions, and unintended uses; (ii) network effects and new flexibility enabled by Grey; and (iii) 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.

2 Introduction

Access control refers to the means by which access to a physical space (e.g., room) or electronic resource (e.g., computer login) is limited to those authorized to gain access. 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 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.

In this paper we present an analysis of a trial deployment of a technology called Grey [2] that is designed to displace 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.

This paper compares the use of Grey to the use of physical keys for accessing offices and other rooms in our university setting. 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 these for

[1] http://www.rsasecurity.com/node.asp?id=1156

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 [6]. 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.

3 Background

In this paper we report on a Grey deployment experiment in which we observed users of physical keys who were given Grey-enabled smartphones to determine the acceptability of this new technology in comparison to the keys it replaced. We begin this background section by describing the use and usability of physical keys. We then provide an overview of the Grey system. Finally, we look at several related studies concerning the usability of access-control mechanisms.

3.1 Physical keys

From a usability perspective, physical keys have a number of positive attributes: they are hard to destroy, work in different environments, and are familiar to the vast majority of users. Unfortunately, they can also be difficult to use for the elderly and those with disabilities [11].

The access-control policies supported by keys are very rigid, allowing little to no change. Once a lock is made, it is difficult to change which keys fit it. Each additional key pattern that fits the lock makes it easier to pick [4], so only a small number of different keys can be made for each lock. To partially combat this issue, many buildings use master and submaster keys. A submaster key usually fits a large number of pre-defined doors, such as all the conference rooms in a building. An administrator can then give a large number of people access to many resources without giving a key to each person for each resource individually. However, this also makes it very difficult to grant someone access to all but one room accessible by a submaster key.

Figure 1: A Nokia N70 displaying the Grey resource list.

Physical keys are also inconvenient to carry, especially in large numbers, and locating the needed key from among several keys on a key chain can be time consuming. In the initial interviews we conducted, participants consistently noted inconveniences associated with keys: large numbers of keys on one key chain can make the correct key difficult to find; and keys are sharp and heavy, making them difficult to carry. Out of 17 users interviewed, 12 had more than one set of keys and 10 left at least one set somewhere when it wasn't in use to avoid carrying more keys than necessary.

3.2 Grey

Grey is a distributed access-control system that uses off-the-shelf smartphones to allow users to access and manage resources. 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.

[2]The wireless phone market is projected to reach 3 billion connections by the end of 2007 [15].

A delegation is represented by a set of one or more digitally signed certificates indicating that a principal (e.g., a named user) has access to a certain resource (e.g., a door) for a specific amount of time. Delegations can be created either proactively before they are needed, or reactively in response to a request from another user. In either case, these certificates are transported via the cellular phone network to the smartphones of the users who need them. Currently, certificates are transfered using SMS messages, while communication between phones and doors occurs over Bluetooth.

To better understand a typical Grey interaction, 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 1. Her phone then contacts an embedded computer governing access to Bob's door.[3] If Bob had proactively delegated the right to enter his office to Alice, then Alice's phone assembles the proper set of credentials and sends them to the embedded computer; if these credentials suffice, then the computer unlocks the door. If Alice did not already have a delegation from Bob, however, then her phone would prompt her to ask someone for a delegation. Suppose Alice selects Bob and sends him a request for access to his office. The request is sent over the cellular network and Bob should receive it wherever he is located as long as his phone is on and receiving a signal. 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. Bob's delegation can be direct or indirect: he could create a new group, add Alice to it, and authorize the group for access to his office. 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.

3.3 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 [16, 1, 8]. 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 are more usable tend to be less secure [5]. 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 [14]. 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 [7]. Other studies showed that low levels of feedback in many access-control systems make it difficult for users to understand what is wrong and what needs to be changed [12, 13]. 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 [13]. These studies look at how users build and manipulate access-control polices. 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 [16, 9]. However, even when users understand how to use an encryption tool, they may fail to use it due to social

[3]In our present deployment, this is a computer embedded in the wall that controls the electric strike on Bob's door. Alice's phone contacts it via Bluetooth.

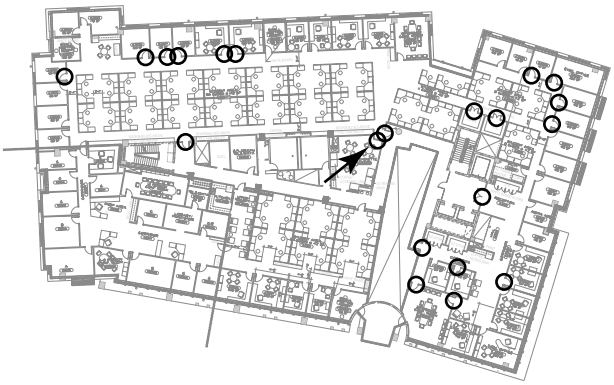

Figure 2: Floor plan of the 2nd floor of the office building. Grey-enabled doors used in this study are circled. Arrow points to kitchenette door.

concerns such as the fear of being perceived as paranoid [10]. The difficulty users have understanding public key cryptography concepts is relevant to our study of Grey, though only tangentially: the concepts of cryptographic keys and encryption are not visible to Grey users. Moreover, 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 [3, 17]. However, as far as we know this is the first published usability study of such a system.

4 Methodology

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.

4.1 Environment

The study was conducted in an office building at a university. Over three dozen doors were outfitted with Grey devices that allowed them to communicate with the phones and to lock or unlock. The Grey phones and doors were set up to log all Grey-related activity.

The second floor of the building (shown in Figure 2) includes a large workspace that is locked after 6pm 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 eleven 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 (not shown in Figure 2). 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.

4.2 Users

In January 2006 we began distributing Nokia N70 smartphones (shown in Figure 1) 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. 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 sixteen 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 sixteen Grey users participating in the study include six computer science and engineering faculty members, six computer science and engineering graduate students, two technical staff members, and two administrative staff members. Fourteen are male and two are female. To preserve privacy we refer to Grey users by fictitious names in this paper.

4.3 Procedure

Before giving a Grey phone to a user we conducted an initial interview that explored their current security practices and how they managed their physical security in the office setting. 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 their use of keys and cell phones.

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 4 to 8 weeks, depending on user availability. 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 changes they made to their access-control policies.

4.4 Videotaping

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 was located near a kitchenette and restrooms that were outside the locked workspace. After hours people regularly used this door to return to the workspace after visiting the kitchenette, restroom, or other areas of the building.

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. The camera was set up to get the widest range possible in the space, capturing an area starting roughly 12 feet from the door. This gave us an adequate range in which to observe users' actions near the door.

4.5 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 3. *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

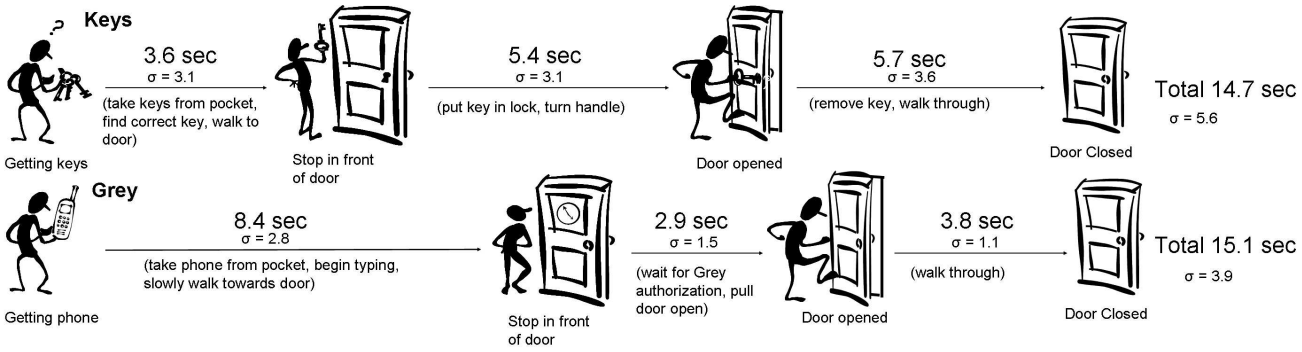

Figure 3: Average door access times for key and Grey users who fetched their key (18 accesses) or phone (7 accesses) in sight of the camera. Most notable is that Grey users spend more time waiting in front of the door but less time moving through it.

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.

5 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 similar systems.

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 variety 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.[4]

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.

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. He apparently found this more convenient than using either his Grey phone or keys to unlock the door when he returned.

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.

[4] `http://fob.webhop.net/`

While to us this seemed relatively fast, we felt it important to understand why our users perceived it as slow. We wanted to observe not only how long Grey took to operate, but also how long 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 [14]. Therefore, we used a video camera to record both key and Grey users accessing a highly trafficked door, as discussed in the Methodology section.

The results of our observations are summarized in Figure 3. 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. 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 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. Our hypothesis is that because Grey users spent a larger fraction of their door-access time simply waiting, it was easier for them than for key users to notice the passage of time and start to feel impatient.

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 17 participants interviewed eight 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 was 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 dialogue 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. A lesson to system designers that we draw from these observations is that they should ensure that perceived performance is no worse than actual performance: perhaps by eliminating steps in which latency will be very obvious to the user, or by making it possible for the user to perform unrelated actions while waiting.

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 the old system. 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.

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. 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 from 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 4, 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.

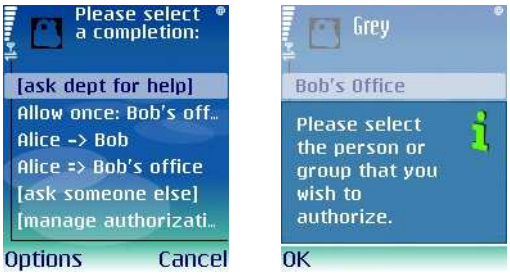

Figure 4: Left: Old reactive delegation interface. Right: New proactive delegation interface.

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. At the same time, though they were generally conscious of the delays incurred by using Grey (e.g., see Principle 1), users did not express any dismay at having to step through wizards in which a screen of instructions followed by an input screen alternated several times. We therefore intend to use wizards for the majority of interactive features that we will add to Grey, and encourage designers of other mobile applications to consider the use of wizards where appropriate.

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.

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 16 users in our study and the 9 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.

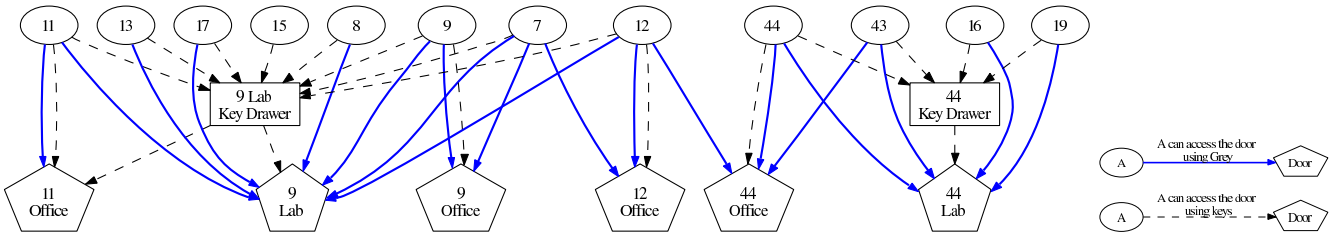

Figure 5: Depicted is a subset of the delegations in use with keys and Grey. Grey users created more direct delegations. Specifically, users were able to avoid the use of key drawers and get direct access to resources.

> 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 have needed one only 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.

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.

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 *key drawers*, hidden public locations where keys are kept for group or emergency use. There were three key drawers maintained by users in our study. In each case the drawer was 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 key drawer or who was using it.

Thomas tells us that the key drawer 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 key drawers. In some cases, people who had access to doors via the key drawers were not issued a corresponding delegation in Grey. In most cases, this was because the users had access via the key drawers as a side-effect of the clumsiness of keys, rather than as part of a desired policy.

For example, with keys and key drawers user 9 had little to no control over who could get into his lab (see Figure 5). As a consequence, user 15, who knew where the key drawer was located, was given inadvertent access to the lab. However, user 9 saw no reason for user 15 to access the lab, so he didn't issue user 15 a corresponding delegation via Grey. Thus, the policy enacted by Grey was more secure (from the standpoint of user 9) than the policy implemented by keys.

With Grey, users also began to delegate access more casually. Of our 16 users, 9 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. As the interfaces of Grey mature and users become accustomed to the technology, we hope to see a larger number of more complex policies (e.g., using groups and roles).

Principle 6: *Unanticipated uses can bolster acceptance.*

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.

6 Discussion

In this paper we have presented six principles related to the adoption of new access control technologies and explained the implications of each for a trial deployment of a smartphone-based distributed access-control system called Grey. The principles themselves are not surprising, and arguably they are widely applicable to many new technologies. However, we did not completely anticipate some of their implications for Grey. The lessons we learned from this field study can inform the development of other access-control technologies as well as smartphone and mobile-device applications in other domains.

Speed and convenience are critical to user satisfaction and acceptance. 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.

A single failure can strongly discourage adoption. 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. Some 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.

Users won't use features they don't understand. Software applications often include a core set of features essential to their operation and additional features that will be used by some users and not others. Users will seek out and use features that perform functions they find useful, but only if they understand how to use them. 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.

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

Low overhead for creating and changing policies encourages policy change. 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. Grey users in our building were able to quickly and easily give short-term or long-term access to other Grey

users. 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.

Unanticipated uses can bolster acceptance. When a new technology is launched, designers often cannot predict the "killer app" that will help it take off. 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.

7 Acknowledgements

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 Grant no. CNS-0627513, U.S. Army Research Office contract no. DAAD19-02-1-0389 ("Perpetually Available and Secure Information Systems") to Carnegie Mellon University's CyLab, National Science Foundation grant no. CNS-0433540, and Office of Naval Research grant no. N00014-04-1-0724.

References

[1] A. Adams and M. A. Sasse. Users are not the enemy. In *Communications of the ACM*, 1999.

[2] 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*, volume 3650, pages 431–445, Sept. 2005.

[3] A. Beaufour and P. Bonnet. Personal servers as digital keys. In *Proc. 2nd IEEE International Conference of Pervasive Computing and Communications*, Mar. 2004.

[4] M. Blaze. Rights amplification in master-keyed mechanical locks. *IEEE Security and Privacy Magazine*, 01(2):24–32, 2003.

[5] 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.

[6] Smart mobile device market growth remains steady at 55%. Canalys Research Release 2006/071, July 2006. Available at `http://www.canalys.com/pr/2006/r2006071.htm` as of Sept. 23, 2006.

[7] X. Cao and L. Iverson. Intentional access management: Making access control usable for end-users. In *Sumposium On Usable Privacy and Security (SOUPS)*, 2006.

[8] 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 Comput.*, 8(6):391–401, 2004.

[9] S. L. Garfinkel and R. C. Miller. Johnny 2: A user test of key continuity management sith s/mime and outlook express. In *SOUPS '05: Proceedings of the 2005 Symposium on Usable privacy and security*, pages 13–24, 2005.

[10] 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.

[11] U. N. Ingrid Thyberg, Ursula Hass and T. Skogh. Survey of the use and effect of assistive devices in patients with early rheumatoid artheritis: A two-year followup of women and men. *Arthritis & Rheumatism (Arthritis Care & Research)*, 51(3):413–421, 2004.

[12] 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.

[13] 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.

[14] U. Piazzalunga, P. Salveneschi, and P. Coffetti. The usability of security devices. In L. F. Cranor and S. Garfinkel, editors, *Privacy and Usability*, pages 221–241. O'Reilly, Sebastopol, CA, 2005.

[15] 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.

[16] A. Whitten and J. D. Tygar. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *8th USENIX Security Symposium*, 1999.

[17] 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.