Non-Stranger Danger

Examining the Effectiveness of Smartphone Locks in Preventing Intrusions by Socially-Close Adversaries

Diogo Marques[1] Tiago Guerreiro[1] Luís Carriço[1] Ivan Beschastnikh[2] Konstantin Beznosov[3]

[1] LaSIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal
[2] NSS, Department of Computer Science, University of British Columbia
[3] LERSSE, Department of Electrical and Computer Engineering, University of British Columbia
[dmarques, tjvg, lmc]@di.fc.ul.pt bestchai@cs.ubc.ca beznosov@ece.ubc.ca

ABSTRACT

Locks are effectively the primary defense against unauthorized access to smartphones. However, unauthorized access by non-strangers seems to be a common occurrence, suggesting that locks often do not fulfill their purpose. To understand the effectiveness of locks, we examine how non-strangers enact intrusions, even when locks are set up. We qualitatively analyzed stories of successful intrusions, recounted by participants in an online study (n = 102). We provide empirical evidence that non-strangers benefit from three inter-dependent factors, that should be taken into account when considering lock effectiveness: trust dynamics, physical proximity, and knowledge of target.

1. INTRODUCTION

Smartphones offer the possibility of setting up authentication barriers which must be overcome at the beginning of interactive sessions – referred to as *locks*. Once locks are surpassed, it is often assumed that user authenticity has been established. Unless there is in-app authentication, device locks are not only the first, but often the last defense against physical interactive access by intruders. It is thus critical that locks are effective.

A consensual approach to improving lock effectiveness is to make locks as convenient as possible, and as strong as possible. Inconvenient locks are thought to be less likely to be adopted by potential users, defeating their purpose; and weak locks are more susceptible to being surpassed by potential intruders. The capabilities of potential intruders, and multiple other factors, however, determine whether locks are either weak or strong. Lock strength can only be defined in relation to a *threat model* that identifies a set of circumstances in which the lock may be tested.

In threat models likely to be relevant to a plurality of smartphone users, one factor that seems important is whether potential intruders are strangers, or whether they are non-strangers (e.g., [2, 14]). Non-strangers, which we define as people within the device owners' closest social circles, such as intimate partners, family, friends, or coworkers, seem better-placed to access devices without permission, if they intend to do so.

Unauthorized access to smartphones by non-strangers does seem to be a common occurrence. A 2012 Pew survey estimated that 12 % US mobile phone owners had at least once experienced another person accessing the contents of their phones in a way that made them feel their privacy was invaded [15]. A 2016 list experiment conducted on Amazon Mechanical Turk (n = 1,381) estimated that, in a 1-year period, 30 % of participants had looked through someone else's phone without permission [11].

It can be reasoned that for intrusions to be so prevalent, existing defenses, such as locks, are insufficiently effective in preventing them. Previous research suggests one reason they may be ineffective is because they are inconvenient, and therefore not used (e.g., [4, 8]). However, even when locks are used, it is unclear whether they are sufficiently strong to prevent non-stranger intrusions.

In this paper, we seek to identify factors that influence lock strength when potential intruders are non-strangers. Non-strangers can be thought of as insiders, in an analogy with large computer systems (e.g., [3, 4]). In such systems, insiders may cause severe damage if not adequately considered in threat models. We similarly seek to inform future smartphone lock threat models that are more realistic and more effective.

To understand the effectiveness of locks in preventing non-stranger intrusions, we examined a set of cases of successful intrusions. We asked participants in a study (n = 102) to recount experiences of unauthorized access by non-strangers, and analyzed cases where locks were mentioned. We found a variety of different sets of circumstances that affected lock effectiveness, which we describe in Section 3. We conclude, in Section 4, by abstracting those circumstances into a set of three interacting factors that differentiate non-stranger threats: *trust dynamics*, *physical proximity*, and *knowledge of target*.

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.
USENIX Symposium on Usable Privacy and Security (SOUPS) 2018.
August 12–14, 2018, Baltimore, MD, USA.

2. METHOD

We recruited participants online, through Prolific, an online survey-taking platform[1]. Being invited for the survey was contingent on participants having previously answered a screening question. The screening question asked whether participants had experienced either someone they knew accessing their smartphone without permission, or themselves accessing a device belonging to someone they knew without permission (we did not ask which of the scenarios was experienced). After obtaining consent, we asked participants to recount those experiences as written stories, through fictional characters "Ash" and "Val". Having previously used this technique to study intrusions of Facebook accounts [20], we found that it stimulates self-reflection, resulting in rich descriptions of incidents.

After excluding nonsensical responses, we kept 102 stories. Participants who provided those stories identified as female in 61 instances, and as male in 40 instances; as being in the 18 to 24 years old range in 31 cases, in the 25-44 range in 63 cases, and in the 45-65 range in 8 cases. We did not enforce limitations on participation based on geography, but information provided by Prolific indicates most participants were located in EU states. The average story was 921 words long (SD = 638) and took 9.4 minutes (SD = 5.5) to write.

We analyzed the collected stories with the help of a piling exercise. To facilitate piling, stories were first printed on cards. Each card was closely read, and text pertaining to the role of locks, if it existed, was underlined to facilitate re-examination. Cards were placed in piles according to the similarity of circumstances in the stories. Descriptive tags were written on sticky notes and placed next to each pile. The piles were iteratively refined, and sometimes subdivided into smaller piles, until we were satisfied that each pile captured a meaningful cluster. Unlike in a coding exercise, we did not develop narrow category descriptions to which new observation could accurately and precisely be assigned to. Our focus was instead on consistence of meaning in each cluster, even if there was some overlap with other clusters. In total, 65 cards were placed in piles. The remaining 37 cards were judged not to reference the role of locks and were excluded in this analysis.

3. RESULTS AND DISCUSSION

We next describe the piles, each composed of stories with similar sets of circumstances. Because our sample is not representative, and this analysis focuses on understanding the diversity of circumstances, and not their prevalence, we do not specify the number of stories in each pile. The illustrative quotes we include were lightly edited for anonymity and clarity. In these quotes *Ash always refers to the person whose phone was accessed without permission*, and *Val always refers to the person who accessed it.*

3.1 No Lock!

While we focus on cases where locks were set but were insufficient to prevent intrusions, we start by examining stories which explicitly referenced locks *not* being set. Prior research has explored the reasons why people do not adopt locks, finding inconvenience to be the most common justification (e.g., [4, 7, 8]). It has also been noted that some people have a perception that their devices do not contain information that would interest others (e.g., [10]). We found examples corroborating both of these explanations, such as:

"Ash had an Android smartphone which was password protected. However, Ash disabled the password protection at some point, because the screen kept timing out when using a GPS program while driving." (P89)

"Val decided to read what Pat was texting to Ash. Ash had never set a passcode on the phone because Ash trusted others and had nothing to hide." (P42)

As the second quote suggests, the belief that trusted people would not perpetrate intrusions, and the belief that untrusted people would not be interested in accessing the device, may together lead users to conclude that there is no value in locking the device. If one of these beliefs is shown to be false, the conclusion may change, as is evidenced in the following example:

"Val accessed Ash's smartphone to frape Ash on Facebook. [. . .] Ash didn't find it funny and put a lock on the phone for security." (P53)

We also observed that a policy of trusting people is sometimes generalized into a policy of trusting locations:

"A quick swipe was all it took to open the contents within. Ash never used password protection at home." (P37)

Adopting such a policy is facilitated by services like Android's SmartLock, which allows setting "trusted" locations where locks are less strictly enforced [6]. Prior research has also suggested reducing authentication requirements and relaxing access restrictions when devices are detected to be in "trusted" locations (e.g., [16, 19]). When considering these approaches, however, it would be beneficial to evaluate their efficacy with a threat model that explicitly accounts for non-stranger intrusions.

We found two other ways in which devices that usually have a lock are accessed without permission at a time when the lock is not active, in ways that seem easier to carry out by non-strangers. One is by anticipating the inactivity timeout, and the other is by borrowing an already unlocked device. Given the specificity of these cases, we divided them into their own piles, which we describe next.

3.2 Beating the Timeout

Locks are usually activated after a period of device inactivity. People known to each other, who are often in the same location, such as those who co-habitate, can take advantage of this known period to attempt an intrusion, like in the following example:

"Ash had left the phone on the kitchen counter of their flat. Ash was only stepping away for a moment to attend to another matter. Val, the flatmate, was able to get access to Ash's phone, as it had been left unlocked, and Ash trusted Val enough to not betray them in this way." (P45)

This type of intrusion is greatly facilitated by the adversary being a non-stranger. Cases where strangers could take advantage of timeouts are also conceivable, but would seemingly require some combination of luck, effort, and motivation. Non-strangers, on the contrary, can enact intrusions of this kind opportunistically, relying on the expectation of

[1] `https://prolific.ac`

trust, and on having multiple opportunities.

Another aspect that may favor these kinds of intrusion is users not being aware of how long it takes for locks to be enforced. There is variation in the length of timeout periods, which can depend on a multitude of factors, such has battery level, OS, or standard preferences. While we found no specific evidence of this, it seems reasonable that non-technical users may not be able to accurately estimate the time to lock activation.

3.3 Borrowing

Another way in which adversaries can take advantage of trust is by having the owner authenticate and lend their device for what seems like an acceptable use. Here's one such example:

"One day, when work was slow, Val asked to borrow Ash's phone. Supposedly, Val wanted to check something on the Internet. Ash unlocked the phone, not thinking twice about the request, and handed it over." (P27)

In some social relationships, refusing a request to borrow a device can constitute an immediate harm to the relationship. Previous research indicates that sharing practices are often related to a desire to communicate trust (e.g., [9, 13]). In such cases, having to go through a lock to share the device may serve a valuable purpose of communicating just how much the other person is trusted.

Strangers can also attempt the strategy of borrowing an unlocked devices, but it seems unlikely that a stranger would be able to use a borrowed device without close supervision.

3.4 Guessing Secrets

One of the ways in which non-strangers can more easily defeat locks is through informed guessing of supposedly secret codes (e.g., [18]).

One kind of information that non-strangers use as candidates for secret codes are dates. We found several references of non-strangers successfully enacting intrusions by guessing dates, such as a birthdate. Here's one example:

"Val opened the phone and was met with a request to enter a PIN number. After several attempts at memorable dates, Val managed to get into Ash's phone." (P48)

This example also highlights that while dates can translate into a PIN in several ways, non-strangers may have several attempts to find the one that works. Guessing secret codes can be used to defeat not only secret-based authentication, but also biometric authentication, such as in the following example:

"Unable to unlock the device without Ash's fingerprint at first, Ash's sibling was able to guess the passcode." (P63)

Biometric unlock methods commonly provide a fallback, secret-based authentication mode. Previous research suggests people do not choose secret codes for fallback authentication that are more difficult to guess [3]. It may even be the case that people choose codes that are easier to remember, fearing that not entering the code constantly makes the code less memorable. Such codes may be easier to guess by non-strangers.

3.5 Shoulder-Surfing

Shoulder-surfing has been a frequent concern for lock authentication (e.g., [7, 8, 17, 22]). Our stories confirm that shoulder-surfing is indeed a vector for intrusions. We found several instances of non-strangers defeating locks by having had previously observed the device owner authenticate. Shoulder-surfing by non-strangers can be particularly effective, given that they may benefit from multiple observations, as evidenced in this quote by P2:

"The phone had a password, but Val, over the last few weeks, had been watching Ash entering it."

Proposals for locks resistant to shoulder-surfing sometimes consider the possibility of extended observation, for instance by allowing a mock attacker to repeatedly observe a video of code-entering (e.g., [1, 21]). However, authentication methods that seem resistant to shoulder-surfing under a single, or very few observations, might not generalize to cases where the intruder is a non-stranger. Recently, Wiese and Roth [22] explicitly modelled shoulder-surfing resistance under repeated partial observation. We found some evidence of intruders engaging in the modelled behavior, including note-taking, for instance in this example:

"A few days later, Ash started noticing Val watching when Ash would type the security PIN to access the phone. Ash also noticed that Val would type something into Val's own phone. Ash's conclusion was that Val was recording the PIN, and then accessing Ash's phone." (P16)

Shoulder-surfing can also defeat the purpose of ostensive security-enhancing practices, such as regularly changing lock codes. Such may be the case especially among non-stranger adversaries, since they may continually shoulder-surf their targets. P75 conveys one example of this behavior:

"Val had watched Ash for some time entering a pattern type password, which Ash changed regularly, but made sure to remember it."

The quotes above illustrate highly motivated perpetrators, who are actively attempting to learn secret codes. In many other instances, however, non-strangers seemed to learn the codes more casually, or even inadvertently.

3.6 Shared and Previously-Known Codes

The most common way we observed non-strangers access devices with locks was with prior knowledge of the access code. Many stories were vague about how the person knew the code, only indicating that the code was known in advance – e.g. *"Val knew Ash's code" (P91), "Luckily, Val knew the password to the phone" (P88), "The phone had a passcode, but Val already knew it" (P35).* In these cases we cannot exclude shoulder-surfing or casual observation. In other cases, however, knowing the other person's code seemed to be tied to the norms of the relationship. Here is one such example:

"Ash and Val were roommates. [...] Val knew Ash's code to unlock the phone, as they both knew each other's." (P34)

Here, there is no indication that codes were shared for a particular purpose. Instead, the proximity of the relationship justifies, by itself, that the codes are known to both parties. The connection between having a close relationship and knowing the other person's code was sometimes explicitly and causally linked, such as in the following quote:

"Since they were friends, Val had Ash's password and took the opportunity." (P77)

It thus seems that, not unlike the dynamic we referenced when discussing borrowing abuse, knowing the code of a close person might serve a social function, such as an indicator of trust. This connection is made explicit, for instance, in the story by P68:

"Val has Ash's access PIN, because they are good and trustworthy friends."

In close relationships, maintaining secrets over access codes could signal as having something to hide, and infringe on norms of openness. We found cases, such as the following, where it was explicitly mentioned that passwords were shared to dispel notions that there could be something to hide:

"Val knew Ash's phone PIN code, since Ash never had anything to hide." (P83)

Aside from communicating trust, people also justify sharing access codes for practical reasons. For instance, they may share codes in anticipation of a future need, such as an emergency [4, 3]. People may also share codes with the expectation that the other party uses the device in a certain, limited way, such as in this example:

"The phone had a code, but Val knew it. Some time before, Ash had told Val the code, in order for Val to make a call." (P21)

We note, however, that even when the impetus for sharing codes is justified with practical reasons, there always seems to be an underlying expectation that the trust deposited in the sharee will be honored. This example also highlights that decisions to share access at one point in time can be difficult to reverse. The fact that a code was shared in the past can be forgotten. If not forgotten, the act of changing a code after it was shared may communicate a lack of trust, which can harm the relationship. To communicate trust, people go so far as to set their devices to unlock with other people's fingerprints:

"Val's fingerprint was set up on Ash's phone: a throwback to the trust and openness that they had once shared together." (P47)

Trust, however, is not a set of stable rules: it is the result of a dynamic process, with considerable potential for asymmetry. At any point in time, it may not be apparent to one person that the other has changed their expectations. The problem is further complicated by the fact that discussing explicit rules of access, enforced only by trust, could itself infringe on trust, and harm the relationship [13].

Explicit sharing of a code can also be difficult to reverse when an intruder escalates temporary access into permanent access. This was described by P44 in a setting with biometric authentication and fallback secret-based authentication:

"Val knows Ash's passcode but has also added their fingerprint to get easy access in case Ash changes their PIN."

Finally, people sometimes can unlock devices because they know other access codes which are re-used as lock codes. Re-using codes is a well-known way to manage multiple demands for authentication secrets (e.g., [3, 4, 8]). Here is one such example:

"Val, being best friends with Ash for a long time, knew Ash's security PIN. Ash used the same PIN for almost everything." (P9)

Stories with code re-use could also have been placed into the "guessing" pile. Our interpretation is that knowing other access codes is distinct from guessing, as it indicates a qualitatively higher degree of knowledge of the device owner's behaviors, which is more exclusive to non-strangers.

3.7 Locks that Worked

Despite almost all the stories describing unauthorized access, in some cases, participants noted cases when locks worked. These cases help improve our understanding of why and when lock are (in)effective.

Unsurprisingly, one type of policy that was successful in some circumstances was setting up a hard-to-guess secret code. In some stories, intruders tried to guess access codes with multiple attempts, but failed. Two of these stories were particularly interesting, since device owners had additional layers of security.

Story 36 notes that *"after enough failed attempts, the smartphone went on a permanent lock"*, which in most configurations means either that a long period of time has to pass until the next attempt, or fallback authentication is engaged. Story 14 relays that not only was the intruder unable to surpass the lock, the device *"had security software that took pictures when someone entered the wrong code"*. A security layer that captures traces of attempted intrusions can reveal unsophisticated intruders. However, more knowledgeable intruders could know of this additional layer, and plausibly deny attempted access by avoiding the camera. More generally, intrusion detection measures could be an effective deterrent: intruders may think twice about attempting unauthorized access if they know that these attempts will be captured and revealed to the device owner. Intrusion detection may also enable device owners to take quick action, as we observed in this example:

"Little did Val know that whenever their phone is accessed without the fingerprint, Ash's smartwatch receives a notification." (P83)

Even when locks can prevent interactive sessions from being initiated, sometimes a lesser degree of access is possible before authentication. We observed some cases, in stories about pranks, where the camera was accessed prior to authentication. We also found cases where intruders inspected notifications available above the lock screen. Although it is possible to set preferences to limit notification appearance or content preview, some users do not set these preferences. And, notifications that could be of little interest to strangers may be highly valuable to non-strangers, including previews of exchanged messages, or notifications (even when lacking content) from sensitive applications, such as dating apps.

4. CONCLUSION

While inspecting the piles of stories we collected, we noticed several common themes. We conclude by describing three factors that we found to explain the ability of non-strangers

to defeat locks; factors that seem distinct from those that would apply to intrusions by strangers.

The first and most critical factor is trust, and, in particular, the **dynamic nature of trust**. Between socially-close people, trust is not only higher than trust between strangers, but it is also more uncertain, as it results from an invisible process. In our stories some people went to great lengths to communicate trust, for instance by lending unlocked devices, or by sharing access codes, or by not obscuring the view when they are entering authentication codes. But, at a later point, this trust was found to be misplaced.

The second factor is **persistent physical proximity**. Non-strangers being habitually close to each other enables most of the intrusions which were recounted. Proximity between parties allows shoulder-surfing, and proximity to devices while they are unattended allows multiple attempts at intrusion.

The third factor that we identified is **knowledge of the target individual**. Knowledge is leveraged in several ways, such as guessing codes or knowing that a code is re-used. But, knowledge also extends to an understanding of an individual's common behaviors, which may allow the intruder to identify optimal opportunities to mount intrusions.

These three factors should inform *realistic* smartphone threat models which explicitly account for non-stranger danger. Some people may find it objectionable to label those in their closest circles as dangerous. However, to many, the most immediate danger to their well-being is indeed a non-stranger – notably, the role of technology in intimate partner abuse is being increasingly recognized (e.g., [5, 12]). Non-strangers are not always dangerous, but lock effectiveness can only be understood and improved by engaging with the full spectrum of experiences faced by users of smartphones.

Acknowledgements

This work was partially funded by FCT - Fundação para a Ciência e a Tecnologia, I.P., through a PhD studentship (SFRH/BD/98527/2013), project mIDR (AAC 02/SAICT/-2017, no. 30347, cofunded by COMPETE/FEDER/FNR), and of the LASIGE Research Unit (UID/CEC/00408/2013).

5. REFERENCES

[1] A. J. Aviv, J. T. Davin, F. Wolf, and R. Kuber. Towards Baselines for Shoulder Surfing on Mobile Authentication. In *ACSAC*, 2017.

[2] A. J. Aviv and R. Kuber. Towards Understanding Connections between Security/Privacy Attitudes and Unlock Authentication. In *USEC*, 2018.

[3] I. Cherapau, I. Muslukhov, N. Asanka, and K. Beznosov. On the Impact of Touch ID on iPhone Passcodes. In *SOUPS*, 2015.

[4] S. Egelman, S. Jain, R. S. Portnoff, K. Liao, S. Consolvo, and D. Wagner. Are You Ready to Lock? Understanding User Motivations for Smartphone Locking Behaviors. In *CCS*, 2014.

[5] D. Freed, J. Palmer, D. Minchala, K. Levy, T. Ristenpart, and N. Dell. "A Stalker's Paradise": How Intimate Partner Abusers Exploit Technology. In *CHI*, 2018.

[6] Google.com. Nexus help: Set up your device for automatic unlock. Last accessed May. 30, 2018. `https://support.google.com/nexus/answer/6093922`.

[7] M. Harbach, A. De Luca, and S. Egelman. The Anatomy of Smartphone Unlocking: A Field Study of Android Lock Screens. In *CHI*, 2016.

[8] M. Harbach, E. V. Zezschwitz, A. Fichtner, A. D. Luca, and M. Smith. It's a Hard Lock Life: A Field Study of Smartphone (Un) Locking Behavior and Risk Perception. In *SOUPS*, 2014.

[9] A. K. Karlson, A. B. Brush, and S. Schechter. Can I borrow your phone?: Understanding concerns when sharing mobile phones. In *CHI*, 2009.

[10] A. Mahfouz, I. Muslukhov, and K. Beznosov. Android users in the wild: Their authentication and usage behavior". *Pervasive and Mobile Computing*, 32:50 – 61, 2016. Mobile Security, Privacy and Forensics.

[11] D. Marques, I. Muslukhov, T. Guerreiro, L. Carriço, and K. Beznosov. Snooping on Mobile Phones: Prevalence and Trends. In *SOUPS*, 2016.

[12] T. Matthews, K. O'Leary, A. Turner, M. Sleeper, J. P. Woelfer, M. Shelton, C. Manthorne, E. F. Churchill, and S. Consolvo. Stories from Survivors: Privacy & Security Practices when Coping with Intimate Partner Abuse. In *CHI*, 2017.

[13] M. L. Mazurek, J. P. Arsenault, J. Bresee, N. Gupta, I. Ion, C. Johns, D. Lee, Y. Liang, J. Olsen, B. Salmon, R. Shay, K. Vaniea, L. Bauer, L. F. Cranor, G. R. Ganger, and M. K. Reiter. Access Control for Home Data Sharing: Attitudes, Needs and Practices. In *CHI*, 2010.

[14] I. Muslukhov, Y. Boshmaf, C. Kuo, J. Lester, and K. Beznosov. Know your enemy: the risk of unauthorized access in smartphones by insiders. In *MobileHCI*, 2013.

[15] Pew Research Center. Privacy and data management on mobile devices. Report. Last accessed May 30, 2018. `http://www.pewinternet.org/2012/09/05/privacy-and-data-management-on-mobile-devices/`, 2012.

[16] O. Riva, C. Qin, K. Strauss, and D. Lymberopoulos. Progressive authentication: deciding when to authenticate on mobile phones. In *USENIX Security*, 2012.

[17] F. Schaub, R. Deyhle, and M. Weber. Password Entry Usability and Shoulder Surfing Susceptibility on Different Smartphone Platforms. In *MUM*, 2012.

[18] S. Schechter, A. J. B. Brush, and S. Egelman. It's No Secret. Measuring the Security and Reliability of Authentication via "Secret" Questions. In *SP (Oakland)*, 2009.

[19] J. Seifert, A. De Luca, B. Conradi, and H. Hussmann. TreasurePhone: Context-Sensitive User Data Protection on Mobile Phones. In *Pervasive*, 2010.

[20] W. A. Usmani, D. Marques, I. Beschastnikh, K. Beznosov, T. Guerreiro, and L. Carriço. Characterizing Social Insider Attacks on Facebook. In *CHI*, 2017.

[21] E. von Zezschwitz, A. De Luca, B. Brunkow, and H. Hussmann. SwiPIN: Fast and secure pin-entry on smartphones. In *CHI*, 2015.

[22] O. Wiese and V. Roth. See You Next Time: A Model for Modern Shoulder Surfers. In *MobileHCI*, 2016.