Zachary Fritchen
CS4960
11/22/15

Progressive Authentication on Mobile Devices

Introduction

Standard authentication schemes on mobile phones are at the moment very limited. They are typically restricted to a single security signal in the form of a PIN, password, or unlock pattern, and the authentication scheme itself is an all-or-nothing affair. This places a greater burden on users than is necessary, which can encourage users to disable security features on their device, rendering their data and applications susceptible to attack. Progressive authentication seeks to increase the convenience of protecting devices with little or no accompanying sacrifice in security.

Forms of Authentication

The various methods of authentication can be grouped into three general categories: what-you-know (knowledge-based), what-you-have (token-based), and what-you-are (biometrics) (Bartik, 2014). Security tokens can be physical objects which the user must carry, and can include car keys, ID or debit cards, USB sticks, and smartphones. Tokens can also be software based, such as cryptographic keys or certificates. Knowledge-based forms of authentication are derived from the user's memory, and include passwords, PINs, and answers to security questions. Biometrics are broken into the subcategories of physiological and behavioral. Physiological biometrics include fingerprint recognition, facial recognition, hand shape or palm veins, and iris recognition. Behavioral biometrics include written signatures, voice recognition, and behavioral patterns.

Zachary Fritchen
CS4960
11/22/15

Multi-factor or multi-level authentication refers to the practice of using more than one signal to determine authenticity. This can increase the level of security, but may or may not decrease the convenience. This is because not all authentication signals place a burden on the user. Web sites may ask for less information when there are secondary signals available, such as IP address or cookies stored on the user's device, or can rely on persistent session cookies which can even eliminate the need for a password to login for a specified period of time (Bonneau et al., 2015).

Modern mobile phones often feature multiple sensors, providing several streams of information which can potentially be used for authentication. These include microphones, cameras, accelerometers and motion sensors, touch sensors, and Bluetooth technology.

Issues with Authentication

The primary issue with authentication is that it exists in a trade-off relationship with convenience. The need to prove your identity repeatedly is a burden to the user, and the greater the security desired, the greater the inconvenience becomes.

Passwords and PINs have become particularly onerous forms of authentication, for multiple reasons. Some of the issues include the greatly varied requirements in length and character type which each authenticating agent requires, the importance of having unique passwords for each system or account the user possesses, and the necessity of some passwords to be changed over time (Hong & Reed, 2013). All of these requirements exacerbate the difficulty of remembering knowledge-based passwords and significantly increases the burden on the user.

Zachary Fritchen
CS4960
11/22/15

Mobile phones provide their own unique problems with authentication. One major issue is that phones are used intermittently throughout the day, and often will deauthenticate themselves after only a few seconds without use. This results in the user having to repeatedly reauthenticate themselves with their device, even if they had just proven their identity seconds earlier. In addition, most security systems on phones are entirely voluntary, giving the user the option of disabling these inconvenient prompts. When faced with frequent authentication prompts, combined with their frustration with existing password schemes, many users will choose convenience over security, putting their data and applications at risk. According to a 2014 survey by Consumer Reports National Research Center, 53% of smartphone users do not use a screen lock on their phone. This fact combined with the increasing prevalence of mobile payment processing or banking applications represents a clear window for abuse. Reducing the burden of authentication on users could potentially lower the barrier of entry for those who have disabled security features on their device.

Standard vs. Progressive Authentication

Many of the issues with mobile phone authentication outlined above are a consequence of the purely binary nature of standard authentication schemes. This “all-or-nothing” method of authentication means the user is either authenticated, and has complete access to the phone, or the user is unauthenticated, and has access to none of the phone (with the exception perhaps of emergency calls and picture taking). This approach does not fit with users’ needs or desires. In a recent survey, users who used security locks wanted about half of their applications unlocked and always accessible, and those with no security locks wanted about half of their applications locked (Riva et al., 2012).

Zachary Fritchen
CS4960
11/22/15

Progressive authentication seeks finer control over the convenience/security trade-off by establishing a non-binary model. Rather than having complete authentication, it establishes levels of confidence in the identity of the user. The necessary security for accessing individual apps or data can be tailored by the user to a specific level based on its sensitivity, effectively reducing the authentication burden for many applications. For example, an app which checks the local weather would require little or no authentication from the user, while access to text messaging history would require a higher degree of confidence.

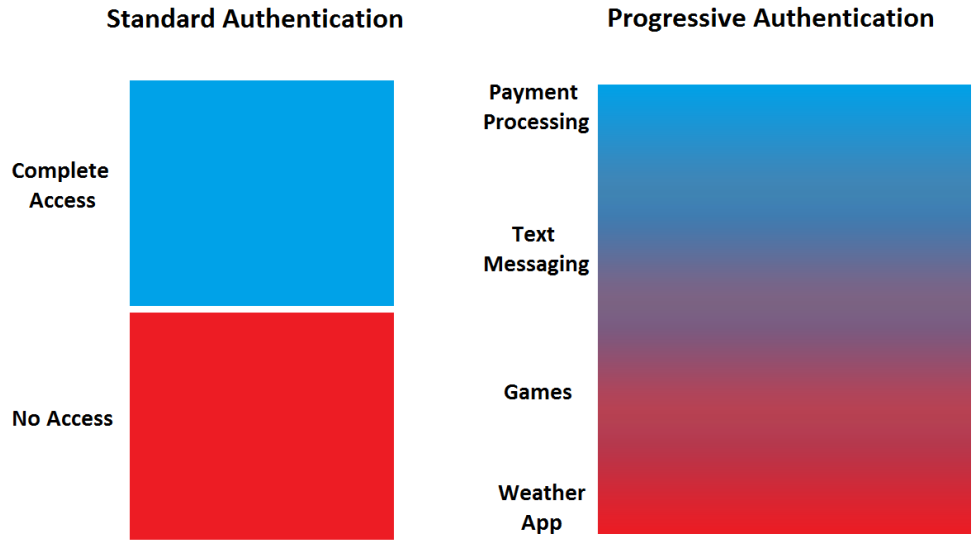

Another fundamental feature of progressive authentication is the introduction of continuity into the authentication model. “Even though the interaction between users and mobile devices may not be continuous, the physical contact between the users and the mobile device can be.” (Riva et al., 2012) A phone which stayed in the user’s possession should maintain its level of authentication over time, reducing the frequency of security prompts. There are various methods and techniques for determining the continuity of possession, including motion and accelerometer sensing, temperature and humidity sensing, voice

recognition, facial recognition, location, and proximity to others devices. Another crucial factor is time, which should generally have a negative effect on confidence as its value grows.

Riva et al. Study Methods

In their 2012 paper, "Progressive authentication: deciding when to authenticate on mobile phones" by Riva et al., an attempt at developing a progressive authentication model was made. Based on a participant survey, researchers decided to utilize three levels of authentication: public, private, and confidential. Public applications required no authentication from the user, private required a medium level of authentication, and confidential required a high level of confidence, usually requiring the input of a PIN. For testing, a Samsung Focus phone was used with Windows Phone 7.1 operating system. A third-party Gadgeteer sensor kit was added to allow the sensing of light, temperature, and humidity.

Nine participants were selected for the study, which was broken into two parts. Part one of the study focused on data collection, overhead such as facial and voice recognition, and machine learning processing. The second part of the study consisted of users following a script designed to test varied use of the mobile device. At multiple intervals attackers attempted to gain access to the phone when it was out of the user's possession.

The following sensors were used to determine authentication level: accelerometers, light, temperature/humidity, touch screen, login events, microphone, and Bluetooth receiver (Riva et al., 2012). A Bluetooth-enabled PC was also placed in the room. Voice recognition was handled by Speaker Sense software, and face recognition was based on a proprietary algorithm.

In order to develop the model for determining authentication level, machine learning with WEKA software was used. Two sets of data were collected: sensory data from the phone,

and observations by the researchers. The observations were used to determine the "ground truth" to serve as objective standards for the machine to learn from. The ground truth was defined according to three labels.

"Public Label: The legitimate owner is not present OR Other people are in contact with the phone OR The legitimate owner is present, but not in contact with the phone and other people are present.

"Private Label: The legitimate owner has been in contact with the phone since the last private-level authentication OR The legitimate owner is present and is not in contact with the phone and no one else is present.

"Confidential Label: The legitimate owner has been in contact with the phone since the last confidential-level authentication." (Riva et al., 2012)

The machine learning model was trained using this data to avoid two specific types of mistakes: false authentication, and false rejection. False authentication refers to the system overestimating the authenticity of the user, and false rejection refers to the system underestimating the authenticity of the user. The former mistake can lead to unauthorized access, and the latter mistake will lead to too frequent requests for authentication. Using the WEKA software, three separate models were produced: decision tree, support vector machine (SVM), and linear regression. Machine learning was also used to model phone placement based on accelerometer and light data. Three placement categories were used: hands, table, pocket.

The figure below illustrates the entire authentication process at work. First, data is gathered and machine learning software is utilized to develop the model. Next, sensor streams from the mobile device and the desktop PC are gathered while the phone is not in use. This is referred to as "low-level processing." Low-level processing also includes the calculation of data such as phone placement, proximity, and voice recognition. When the touch screen is activated, the "high-level processing" takes place, which includes extracting the feature vectors from the

data and plugging them into the model to determine the authentication level. These high-level calculations, as well as voice identification, can be offloaded to a remote device or the cloud in order to reduce the power consumption and performance load on the mobile phone.

A risk factor (R) was also implemented, which allowed the researchers to regulate how aggressive the models were in authenticating users. This value ranged from 0.05 to 20, with the smaller values representing lower levels of risk and thus more aggressive security. These risk values could potentially be configured by the user of the device to choose their own preferred level of security, thus giving finer control over the convenience/security trade-off.

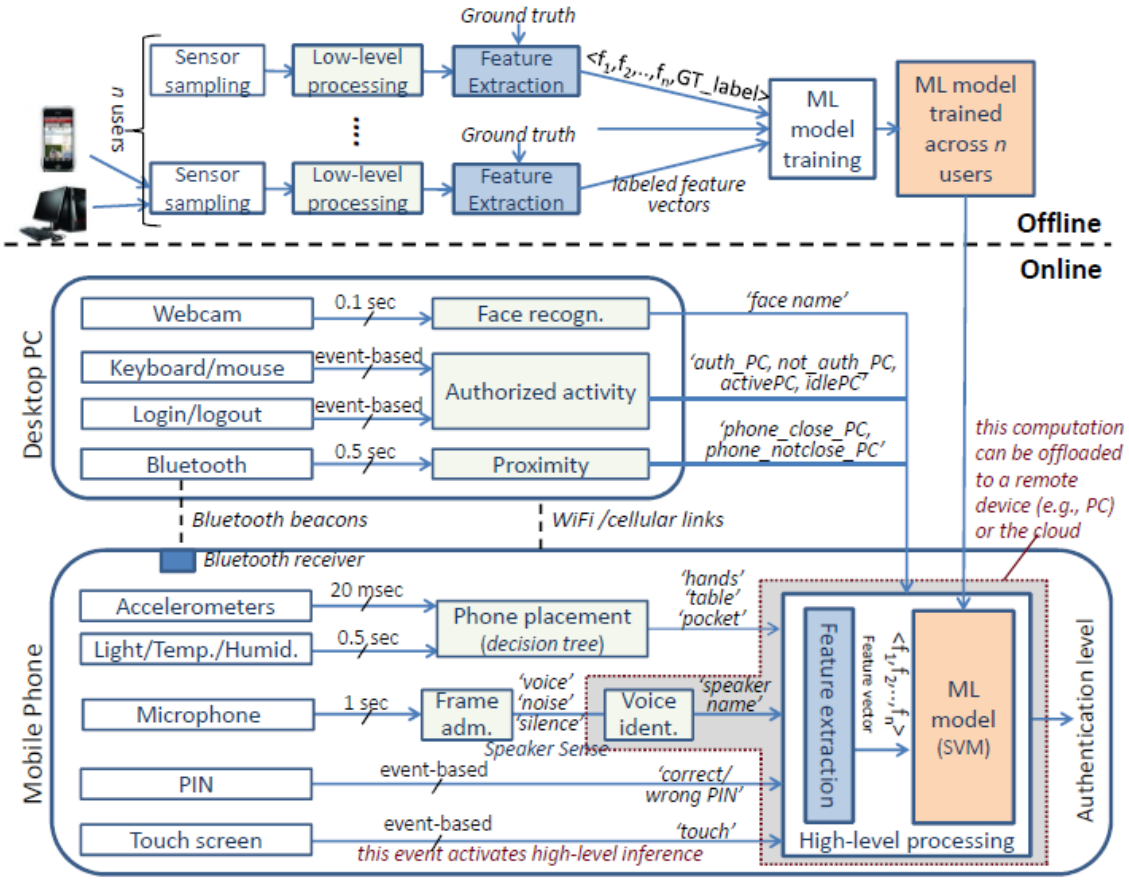

(Riva et al., 2012)

The table below lists all of the variables used in the high-level processing, as well as a description of their meaning. The majority of the variables store times since a logged event, in order to determine the continuity of possession. The features which had the strongest correlation with ground truth were ProxAuthDev (proximity of authenticated device), LastPlacementDuration, TimeSincePin, and TimeSinceOwnerVoice, in that order.

Table 1: Machine learning features used in the high-level processing.

Category	Features	Description
Cont.	Placement, PlacementDuration, PlacementConf	Current placement of the phone, how long it has lasted, and associated confidence
Cont.	LastPlacement, LastPlacementDuration	Last placement of the phone, and how long it lasted
Cont.	TimeSinceTable, TimeSinceHands, TimeSincePocket	Time elapsed since the last time the phone was on the table, in the user's hands, or pocket
Cont./Secrets	TimeSincePIN, TimeSinceTouch	Time since last login event and time since the phone's screen was last touched
Biom.	Speaker, SpeakerConf	Whether a human voice was identified (owner, other, no-voice) and associated confidence
Biom.	TimeSinceOwnerVoice, TimeSinceNonOwnerVoice	Time since (any) voice was identified
Biom.	TimeSinceSound	Time since any sound (either voice or noise) was detected
Poss./Biom.	ProxAuthDev, ProxAuthDevConf	Proximity of phone to a device where the user is logged in and active, and confidence
Poss./Biom.	TimeSinceProx	Time elapsed since the proximity status last changed

(Riva et al., 2012)

Riva et al. Study Results

The participants using phones with standard authentication were required to input their PIN an average of 19.2 times, while the progressive authentication group required an average of 11.2 entries. This represents a 42% decrease in authentication requests for the user. Both groups experienced 0.0% unauthorized authentications (UAs), which refers to a non-legitimate user attempting to access the phone. Therefore, the progressive authentication model resulted in increased convenience with no resulting loss in security.

The rate of required authentications for Public applications was reduced by 100% because the Public setting required no authentication. The rate of required authentications for

Confidential applications was reduced by 0%, because these applications require the highest necessary confidence. The results for Private applications depended on the risk factor value being applied, which ranged from low-risk (0.05) to high-risk (20). The rate of false rejections (FRs) for Private applications ranged from 57.7% for low-risk levels to 34.4% for high-risk levels. This means that users were required to authenticate to access applications with Private security level roughly 58% of the time for aggressive security settings and only 34% of the time for more passive security settings. The rate of false authentications (FAs) for Private applications ranged from 3.3% for low-risk to 16.1% for high-risk.

The model accuracy varied from 83% to 100% when all sensors were used. The least accurate sensor used was voice identification, which only properly recognized the owner's voice 77% of the time. There was also a bit of inaccuracy with the placement sensors, which confused a pocket for a table roughly 6% of the time. This was likely due to the phones light sensor sticking out of the user's pocket and thus detecting light when none was expected. The most accurate was face recognition with 94% accuracy and little variance across users. All of these represent areas where the progressive authentication scheme could be improved to reduce the occurrence of false rejections or false authentications.

Of the three models tested, the support vector machine (SVM) showed the best results and had the highest precision and recall, while linear regression performed the worst. Precision is defined as "the fraction of correct predictions across all testing samples that resulted in the same prediction." Recall is defined as "the fraction of correct predictions across all testing samples with the same ground truth label." Linear regression experienced many incorrect predictions and was rejected by the team. The decision tree model was more aggressive and

experienced fewer false rejections, but generated more false authentications. SVM had very few false authentications and a high precision for all labels.

Drawbacks of Progressive Authentication

There are four primary drawbacks of progressive authentication: increased power consumption, impact on performance, initial overhead, and privacy concerns. The low-level processing in which phone sensors are constantly active and collecting data increase the power consumption on the phone. The researchers experimented with four separate power configurations in order to combat this problem. For comparison, the idle power consumption of the phone, in which only the screen and WiFi were active, was 896 mW. The *Local* configuration, in which all processing and calculations take place on the phone, had the highest power consumption at 651 mW. However, by off-loading some of the power intensive calculations to a cloud service or secondary PC, they were able to reduce power consumption significantly. The *Remote* configuration, in which both low and high-level processing were outsourced, experienced a power consumption of 307 mW. The best rate of power consumption was achieved by the *LocalMin* configuration, which disabled the high-power consuming tasks of Bluetooth proximity detection and voice identification, and resulted in only 42 mW of power.

The low and high-level processing are also quite calculation intensive. The model must calculate the authenticity level based on all the factors involved, and those factors must also be calculated, such as identifying faces from images or identifying phone placement from accelerometer data. These calculations take time, and thus increase the latency of the system. The *Remote* configuration experienced the worst performance, with 1.5-2.81 seconds of delay,

likely due to the data transmissions required. The *Local* configuration had an execution time of 0.23 seconds.

The best compromise was found in a *LocalRemote* configuration, in which computation-light and low-level processing are executed on the phone, while calculation intensive and high-level processing such as voice identification are off-loaded. This configuration resulted in 325 mW of power consumption and 0.99 seconds execution time. Having a range of configurations on the phone can improve performance and power consumption even more by switching between configurations and in some cases temporarily disabling progressive authentication features.

Table 8: Power consumption and execution time on a Samsung Focus WP 7.1 for 4 different power configurations.

Conf	Sensing (mW)	Comput (mW)	Comm (mW)	TotalPower (mW)	ExTime (sec)
LocalMin	<1	41	0	42	0.20
Local	≈160	447	44	651	0.23
LocalRemote	≈160	71	94	325	0.99
Remote	≈160	49	98	307	1.50-2.81

(Rita et al., 2012)

The progressive authentication model used also suffered from some initial overheard. This includes time spent establishing face and voice recognition, time spent setting application security level, and the registering of trusted devices. The researchers suggested that voice recognition overhead can be reduced by collecting the necessary data during normal phone conversations.

A fourth issue, which was not addressed in the study, is privacy concerns. The progressive authentication model described was constantly collecting streams of data from the user, such as proximity to other devices, the voice of the owner or others in the room, and the faces of the owner or others in the room. The models used would also be required to permanently store some of this data. It is possible that companies which implement this type of software on the phones, or law enforcement or intelligence agencies which can pressure such companies to cooperate, may engage in data mining which would put users' privacy at risk. It is also possible to imagine attackers gaining access to these data streams and collecting personal information from users.

Conclusion

Progressive authentication has been shown to be a viable means to improve the convenience of authentication methods on mobile phones without sacrificing security. It carries the drawbacks of increased power consumption, decreased performance, and initial overhead. Additional studies and more refined models may help to reduce these drawbacks over time. Reducing the frequency of security prompts is an important step towards encouraging mobile phone users to protect their data and applications.

Zachary Fritchen
CS4960
11/22/15

References

Bartik, C. (2014, January 10). 3 different types of user authentication. Retrieved from http://www.cloudentr.com/latest-resources/industry-news/2014/1/10/3-different-types-of-user-authentication

Bonneau, J., Herley, C., van Oorschot, P. C., & Stajano, F. (2015). Passwords and the evolution of imperfect authentication. *Communications of the ACM*,*58*(7), 78-87.

Consumer Reports National Research Center. (2014). Annual State of the Net Survey. Retrieved from http://www.consumerreports.org/cro/news/2014/04/smart-phone-thefts-rose-to-3-1-million-last-year/index.htm

Hong, J., & Reed, D. (2013). Passwords Getting Painful, Computing Still Blissful. *Commun. ACM*, *56*(3), 10-11.

Riva, O., Qin, C., Strauss, K., & Lymberopoulos, D. (2012, August). Progressive Authentication: Deciding When to Authenticate on Mobile Phones. In *USENIX Security Symposium* (pp. 301-316).