

Securing Smartphones: A Micro-TCB Approach

Yossi Gilad¹, Amir Herzberg¹, and Ari Trachtenberg²

¹Computer Science Department, Bar-Ilan University

²Electrical and Computer Engineering Department, Boston University

Abstract

As mobile phones have evolved into ‘smartphones’, with complex operating systems running third-party software, they have become increasingly vulnerable to malicious applications (malware). We introduce a new design for mitigating malware attacks against smartphone users, based on a small trusted computing base module, denoted μ TCB. The μ TCB manages sensitive data and sensors, and provides core services to applications, independently of the operating system. The user invokes μ TCB using a simple *secure attention key*, which is pressed in order to validate physical possession of the device and authorize a sensitive action; this protects private information even if the device is infected with malware. We present a proof-of-concept implementation of μ TCB based on ARM’s TrustZone, a secure execution environment increasingly found in smartphones, and evaluate our implementation using simulations.

Keywords: Security Kernels; Invasive Software; Smartphones; Trusted Physical Interfaces

A version of this paper will appear as:

- Y. Gilad, A. Herzberg, and A. Trachtenberg, *Securing Smartphones: A Micro-TCB Approach*, IEEE Pervasive Computing Magazine, 2014.

1 Introduction

The increasing popularity of cellular phones has steadily shifted personalized computing from wired desktop machines to wireless mobile devices. Accompanying this shift was a natural move to utilize mobile phones as ‘trusted digital agents’ [1], allowing users to receive private messages and perform mobile payments or other confidential tasks. Early mobile phones were simple, limited, and had proprietary implementations, leading to the impression that they were more secure than their multi-purpose desktop counterparts, on which users could install unchecked, and possibly malicious, third-party software.

Today’s cellular phones, however, have evolved into *smartphones*: general purpose, complex, computer/sensor devices that employ multiple wireless communication technologies and are ever connected to both the Internet and their users. In effect, today’s smartphones are a natural extension of a person’s daily life: texting/emailing friends, catching up on news/stocks, and managing bank accounts. Not only do they store passwords, account numbers, pictures, and videos that were freely saved on yesterday’s phones, but they also place numerous sensors well within the comfort zone of their users.

It is therefore not surprising that the evolution to smartphone devices arrived with a concomitant rise in exploits and attacks. The complex operating systems and variety of applications in today’s smartphones

pose a rich attack surface that was not present in the earlier phones. Malicious applications steal user data, gain administrative privileges by exploiting software vulnerabilities [2] and steal user credentials by faking legitimate notifications [3]. Worse yet, the vulnerability/patch cycle for mobile devices is often long, due, in part, to the potential need to modify the code of several phone manufacturers [4]. Finally, many smartphone users are non-technical, making security an even greater challenge.

A successful security architecture must rely as little as possible on the user’s skill-set. It should be simple and small enough that its execution paths can be evaluated for security, and yet sufficiently elemental to support a variety of services. In this paper we propose one such security architecture, which uses the phone’s physical proximity to its user as a basis for authentication.

Threat Model. This work considers attacks by malware applications. The standard defense against such attacks is *application sandboxing*, wherein a permission system offers applications a variety of approved privileges to resources, ranging from sensor access to the ability to display notifications. These privilege-restrictions limit the potential damage of malicious or exploitable applications.

However, the privilege systems fail if an application ‘breaks out of jail’ by exploiting vulnerabilities in the (complex) smartphone operating system or a privileged third-party application. Indeed, even security-focused application-level defenses (e.g., [5]) may fail when the operating system itself is compromised by malware. Under such conditions, privileged malware can read files, monitor the keyboard, record the microphone, take pictures and learn the user’s location, all posing a significant threat to users’ privacy.

In order to function even under operating system compromise, our security design is based on a *minimal* Trusted Computing Base (μ TCB) that supports a few necessary operations. We assume that the μ TCB implementation executes in a dedicated environment and is small enough to be thoroughly validated, we thus preclude attacks on this component. Our security design further assumes that the user currently possesses the phone, and thus we preclude physical attacks.

1.1 Design Overview

We integrate a simple and minimal trusted component into the smartphone, and migrate key security-oriented services from the base operating system (Android, iOS, etc.) to this component. In a way, this component resembles the micro-kernel approach to secure operating system (OS) design, in which the OS is reduced to a minimal kernel and most functionalities are implemented outside the kernel; we thus name our trusted component a *Micro Trusted Computing Base* (μ TCB).

The core of our design is the separation between the smartphone’s complex OS and the μ TCB, which allows applications that execute over the smartphone’s OS to access its services using a programming interface. In addition, when the user presses a dedicated button, called the *Secure Attention Key (SAK)*, the device enters a special *secure-mode*. In secure-mode, sensor signals and hardware interrupts are first processed within the μ TCB environment, and only then possibly relayed to the OS. This allows μ TCB to take control of peripherals, such as the screen, and interact with the user securely even if the OS or other phone subsystems are compromised. By pressing the SAK, the user validates that she physically possesses the device, thereby distinguishing between automated or remote instructions and user commands.

We demonstrate the applicability of our design with two fundamental examples, secure messaging and mobile-payment applications; both execute within the smartphone’s OS but base their security on μ TCB services. As a proof-of-concept, we implement our μ TCB design within the ARM TrustZone security environment and provide a performance evaluation using ARM’s device simulator.

2 Related Works

2.1 Trusted Physical Interfaces

A secure path between the user and a trusted system component is a key requirement of a usable security architecture, and this is particularly true for smartphones [6], which are widely used by naive users. The concept of using a dedicated secure attention key (SAK) to establish such a path was introduced in [7] for opening a secure login-shell. This design is incorporated in modern operating systems, for example, Windows’ Ctrl+Alt+Del combination, but is typically limited to login or task management, and can be affected by a compromised operating system.

Our design adapts this idea of a dedicated button to smartphones. We secure private user data, such as login passwords, files, and messages, and enable secure management of sensitive sensors. Finally, we provide a programming interface that allows complex (external) applications to delegate their security building-blocks to the simple μ TCB core.

Among a variety of related works, McCune et al. considered the problem of user-observable path verification [8], where the user receives a clear and unspoofable feedback that she communicates securely with a trusted endpoint. To this aim, [8] suggests using a LED indicator connected to the path endpoint. Our design includes a similar indicator, attached to the μ TCB, which signals when the device is in secure-mode and the user can interact securely with the μ TCB. However, we do not solely rely on passive indicators, combining additional mechanisms to create a ‘defense in depth’-like architecture: we train the user to press the SAK prior to every sensitive action (ensuring that the device is in secure-mode), as well as force the user to enter secure-mode to perform certain sensitive operations (such as signing documents).

2.2 Trusted Computing Base (TCB)

TrustVisor [9] is a small TCB, designed for personal-computers, that allows applications to execute security-sensitive codeblocks (e.g., signature verification). Complementarily, the Trusted Language Runtime (TLR) [10] architecture provides similar capabilities for smartphones. While TrustVisor and TLR provide important security guarantees, they are strictly designed to secure *computations* and not interactions with the user; for example, a malware that replaces the keyboard device driver could, under both systems, record or modify the user’s keystrokes.

TrustVisor has been extended to create a trusted path between an input/output (I/O) device and an application [11]. The design of [11] requires that each application implement its own trusted I/O device functionality (as part of the TCB); in contrast, our μ TCB provides trusted services to untrusted applications (through a restricted programming interface). This makes it easier to code applications, but also makes our design less flexible (e.g., only the manufacturer can introduce new TCB primitives). We believe that codability is essential for today’s constantly growing smartphone application markets: only the most sophisticated developers can be expected to properly write and *thoroughly verify* their TCB application code. In addition, the verification of the trusted path in [11] makes use of an additional hand-held device, which can be an impediment to consumer adoption; our design *embeds* a trusted path between the user and μ TCB, which the user can enable by pressing the SAK. This allows us to incorporate the verification indicator into the smartphone, by placing it under the control of the μ TCB.

The trusted sensors architecture of [12] uses a TCB to provide reliable sensor readings for applications, and Raj et al. [13] show how a TCB could monitor network traffic and attest for its volume. These architectures provide important services, but do not intrinsically prevent malware applications from obtaining sensor readings or accessing sensitive files; they are thus complementary to our design.




Figure 1: The μ TCB Architecture, Components and Trust Assumptions.

3 The μ TCB Design

In this section we present the μ TCB architecture and its interaction methods with the user.

3.1 Architecture

The μ TCB platform has two important properties. First, it provides a code-execution environment that is separated from the smartphone’s OS, allowing μ TCB to remain secure even when the OS is compromised. Second, μ TCB code running over this platform can handle or filter signals from sensors and peripherals *before* they are handled by the OS. This property allows privileged access to the secure attention key and touch-screen signals, and enables μ TCB services to confidentially interact with the user (see next subsection).

The μ TCB is minimal by design (in order to avoid vulnerabilities), and it thus relegates all but a few security-oriented services to the smartphone OS; these services, described in Section 4, are accessible to applications running on smartphone’s OS via a restricted API. Though minimal, μ TCB services allow implementation of complex and fundamental applications, as we show in Section 5.

Figure 1 illustrates our architecture at a high level and the system’s ‘trust assumptions’ in participating entities, implementation details will be covered in Section 6.

3.2 The Secure Mode and Attention Key

The μ TCB has a dedicated Secure Attention Key (SAK) that is not visible to the smartphone’s OS. When the user presses the SAK, she triggers an interrupt that invokes an interaction with the μ TCB system; we call this interactive mode *secure-mode*. The μ TCB presents to the user a menu of simple services as well as pending requests from applications, we describe those services and requests in Section 4.

Pressing the SAK *ensures* that the smartphone is in secure-mode, namely, the SAK is a ‘one-way button’. In order to exit the secure-mode, the user presses a virtual ‘Exit’ button in the secure-mode menu. The smartphone also automatically exits secure-mode after idle timeout (e.g., five minutes), this is to avoid confusing the user regarding the smartphone’s mode after a long idle period.

When in secure-mode, μ TCB temporarily handles signals of the touch-screen and filters those of other sensors, blocking them from reaching the smartphone’s OS so as to ensure that no application, not even one that escapes ‘out-of-sandbox’, may view or manipulate the screen to leak private user inputs.

Secure-Mode Indicator. A key requirement of our system is that users receive a clear and unspoofable indication when their smartphone is in secure-mode. We identify two cases where such indication is crucial: (1) the user enters secret data, (2) the user disables a sensor. In the first case, an application that tricks the user will learn the secret data; in the second case, a sensor that remains active might expose sensitive

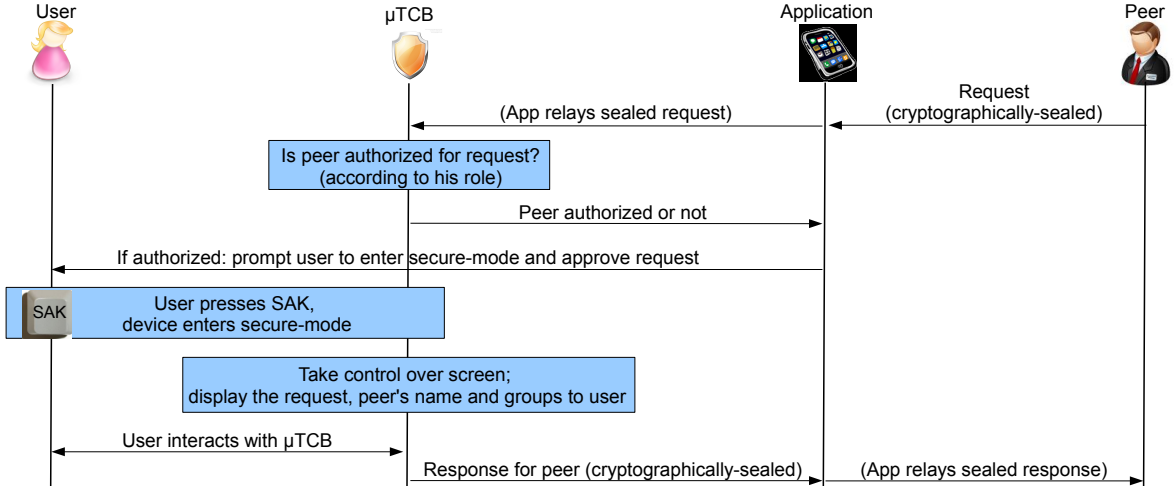


Figure 2: Role-Based Access Control API Scheme.

information. To this aim, a LED indicator is connected to the μ TCB and lights only when the device is in secure-mode (the LED is inaccessible to the smartphone’s OS).

Forcing and Training Ceremonies. Usability studies suggest that users often do not notice passive security indicators; to improve security, we use *forcing and training ceremonies* [14, 15].

First, the user’s private data and keys are only accessible to μ TCB, *forcing* the user [14] to press the SAK (and enter the ‘real’ secure-mode) in order to use them, e.g., to sign documents with her private key.

Second, we support a *training mechanism*, shown to improve indicator awareness [15]. We assume an initial period where the smartphone is infection-free, and simulate a spoofing attack by occasionally neglecting to light the indicator LED when the smartphone enters secure-mode. Any action by the user without repressing the SAK (i.e., to ensure that the device is in secure-mode and indicator turns on) results in a negative feedback, e.g., the LED blinks and μ TCB notifies the user to notice the indicator.

4 Services and Interfaces

The μ TCB provides security-oriented services to the user as well as to applications running on the smartphone’s operating system. These services are available through restricted user and application interfaces.

Our application programming interface (API) embeds *role-based* access restrictions, as illustrated in Figure 2, to allow applications to access μ TCB services only *on behalf of* a remote peer (representing a remote entity like a cloud service, contact person or online store) that has permission to use the service. In such cases, μ TCB uses a public key assigned to the peer to cryptographically seal/unseal data for/from the calling application; in essence, role-based access control allows the application to only relay, but not read or modify, the sensitive information and, therefore, prevents exposure of such information in the case that the application is malicious.

We next present the services and interfaces that μ TCB offers, organized into four categories: authorization, confidentiality, non-repudiation and sensor control.

4.1 Authorization of Peers, Roles and Groups

The μ TCB restricts sensitive operations, such as those that disclose private information, to authorized peers. Each peer is represented by a name assigned to a public key. The peer is authorized for permissions in two logically orthogonal classes: *groups* and *roles*.

Groups restrict the distribution of information, and can be modified or created (by a certificate authority or the user); for example, work-related documents may be restricted so that they are only sent (under encryption) to the group of employees of the company.

Roles, on the other hand, restrict the actions permitted by peers. There are exactly three roles:

Contact: this role enables a peer to send/receive encrypted and authenticated information to/from the user.

Signatory: peers with this role can send signed documents and request a counter-signature from the user.

Certificate authority (CA): peers with this role are trusted to identify and authorize other peers. A CA is restricted to authorizing peers of specific groups (e.g., a company’s CA can only authorize its employees). Credentials of well-known CAs are installed to μ TCB before user-adoption, e.g., by the manufacturer.

User Interface and API. The user may authorize new peers. However, this management option is risky and should typically be used only by system administrators to authorize a local CA; as such, this option may be ‘locked’ with an administrative password.

The authorization service is inaccessible to applications, to avoid exposing the user to social engineering attacks (e.g., where the malware authorizes a peer with similar name to the user’s bank).

4.2 Confidentiality

Smartphones often store sensitive data, ranging from login credentials, to private pictures and messages; μ TCB provides means of ensuring the confidentiality, authenticity and integrity of such data, and an infrastructure for distributing it among *contact* peers.

More precisely, μ TCB includes a *private repository*, which is essentially a secure file system (inaccessible to the OS). Some of the user’s files in the repository, and in particular her *private key*, are encrypted with a user-provided credential such as a password; the user’s credential is never saved, so that she must be directly involved to access this information, helping to mitigate the risk of identity theft from a stolen phone.

The μ TCB can also encrypt and authenticate data for transmission to contact peers. In this case, the sender’s μ TCB uses the recipient’s public key to encrypt two symmetric keys: one to encrypt and the other to authenticate the information. The encrypted keys are signed with the sender’s private key, which allows the recipient to validate the authenticity and integrity of messages from the sender. We encapsulate messages using the “encrypt-then-MAC” paradigm, which permits verification of the message’s authenticity prior to decryption, to quickly discard spoofed messages.

User Interface and API. The user can access the private repository as a general purpose file system. Each file is associated with an access list of contact peers, identified by their names or groups, who may receive it.

Applications can use the μ TCB to handle private information on behalf of authorized *contact* peers (following Figure 2). The API functions that access the confidentiality service state the intended recipient or the sender of the data, and μ TCB validates that he is an authorized contact; the validation result is returned to the application. If the peer is authorized, then the application prompts the user to press the SAK and enter secure-mode to interact with μ TCB. The user can view the peer’s name and groups while in secure-mode. The two confidentiality-service API functions operate as follows:

Request-Data(*recipient*). The μ TCB requests the user to enter a message, select a file from the private repository (that *recipient* is authorized to read) or send sensor-readings (e.g., camera picture, location coordinates); μ TCB then encapsulates the information as described above and returns the cryptographically-sealed data to the application.

Display-Message(*sender, message*). The μ TCB displays *message* along with *sender*'s name and groups.

4.3 Non-Repudiation

The μ TCB allows the user to sign documents presented by *signatory* peers, possibly after the user completes her personal details (e.g., name, identification number, etc.), while ensuring that neither the user's private key nor her personal details leak to OS applications.

Each document has a type (e.g., banking, commerce, etc.) and some important types are reserved and may be presented only by peers in specific groups (e.g., only a bank may present a banking form).

User Interface and API. The μ TCB stores the documents that the user has signed, and allows the user to later view those documents. Applications can request the user to sign or display documents on behalf of *signatory* peers via the following API functions (that implement the scheme in Figure 2):

Request-Signature(*recipient, document*). When an application calls this function, μ TCB validates that *recipient* is an authorized signatory peer, that *document* is signed by *recipient*, and that *recipient* may request signature for this document type.

When the user presses the SAK and places the smartphone in secure-mode, she can view the document and possibly complete required personal details; if she approves the signature request, then μ TCB uses her private key to counter-sign the document (the user signs the completed document and the original signature of *recipient*). The application receives the user's signature and personal details encrypted with *recipient*'s public key.

Display-Signed-Doc(*sender, document*). This function is similar to the previous one, but only uses μ TCB to display *sender*'s signed document to the user (i.e., the user does not counter-sign it).

4.4 Sensor Control

The sensor-rich smartphone allows privileged applications to record, take pictures and locate the user *at any time*. Existing smartphone operating systems typically permit users to turn off some sensors (such as the GPS) for power conservation, but other sensors (such as the microphone or camera) are not directly controllable. Moreover, on an infected device any of these sensors may be activated without indication to the user.

In our architecture, the sensors signals are routed to the OS via μ TCB, allowing it to completely block their signals from the OS.

User Interface and API. The user can block or activate the sensors. Applications can request the user to *temporarily* enable a sensor via the **Enable-Sensor(*sensor*)** function. When an application calls this function, μ TCB validates that *sensor* is currently blocked (or will be re-blocked soon, see below).

When the user presses the SAK, she can approve activating the sensor for a short period of time (e.g., up to an hour), or chose to automatically discard activation requests for some time period. The μ TCB notifies the application of the user's decision and its timeout, allowing the application to again request to activate the sensor when the decision expires.




Figure 3: Secure Messaging. User interface and training feedback (in virtual machine).

5 Applications

In this section we describe example applications that illustrate how μ TCB and its API allow implementation of key security-demanding applications.

5.1 Messaging

The secure-messaging application includes a check-box where the user can mark messages as ‘secure’, see Figure 3 (a). In order to send a secure message, the application retrieves the recipient’s certificate (which includes his public key, groups and role) by requesting it from the recipient’s device, e.g., via SMS or network interface (the certificate is signed by a CA, and signature validation performed by μ TCB).

When the user sends a secure message, the application calls the Request-Data API function, and prompts the user to press the SAK and enter a message for that recipient, as in Figure 3 (b). The μ TCB returns to the application the *cryptographically-sealed* message, and the application then sends the message to the recipient.

When the application receives an encapsulated message, it calls the Display-Message API function and prompts the user to enter secure-mode to read the message.

In order to encourage users to incorporate the secure functionality, the secure option is set as default for communication with contacts to whom the user had sent/received a secure message before. If the user changes this configuration and attempts to send a plaintext message to such contact, the application alerts the user, as in Figure 3 (c).

5.2 Commerce and Payments

We next present a mobile payment protocol that integrates the μ TCB to mitigate the risks of identity theft and exposure of payment credentials. Our description uses the example of a stock trading application for sake of exposition. We assume that prior to running the protocol, the user has registered her public key with her bank.

Protocol. In the *offer* phase, the user requests to buy a particular stock using a smartphone application. The application might find several available brokers, all of which should be *signatory* peers approved by trusted CAs (e.g., banks), and it presents to the user the information required to decide on a particular offer (e.g., the broker’s commission). The application calls the Request-Signature API function, which provides to μ TCB the identity of the selected broker and his signed offer.

In the *payment* phase, the user enters secure-mode and views the broker’s offer; if she decides to accept it, then she enters the secret credential (e.g., password) for decrypting her private key file in the repository (see Section 4.2). The μ TCB uses the private key to sign an order after the user completes relevant personal details (e.g., account number). The signed order is cryptographically sealed using the broker’s public key and the cipher-text is returned to the application which then relays it to the broker.

Finally, in the *confirmation* phase, the broker decrypts the user’s response and clears the payment with her bank. The broker sends a signed receipt to the application, which then calls the `Display-Signed-Doc` API function and notifies the user to enter secure-mode to view the receipt. If the user does not receive a receipt in a timely fashion, she can notify her bank and revoke the transaction (this may be automated).

Security Benefits. This payment protocol improves typical payment applications since it mitigates the risk for identity theft: the user does not expose her payment credential (i.e., private key) to the bank, broker or the application (cf. to credit-card payments). Moreover, even if her smartphone is lost or stolen, a user-chosen credential is required to decrypt her private key, further reducing the risk for impersonation or fraud.

6 Implementation

In this section we describe our Proof of Concept (PoC) implementation of μ TCB over ARM TrustZone, a trusted code execution platform already available in some of today’s mobile devices. TrustZone separates between the *normal-world* (*NWorld*), where the rich OS and applications execute, and the *secure-world* (*SWorld*), where the μ TCB resides.

We implemented a *proof of concept* (PoC) prototype to verify that implementation of our design over TrustZone is feasible. Our PoC (available at <https://sourceforge.net/p/cellsec/source>) implements the μ TCB services and provides an API (in form of a library) for applications running in *NWorld*; our TCB has less than 8000 lines of source code and is much smaller than the phone’s OS kernel. We note that our PoC does not modify the smartphone’s OS to provide applications access to TrustZone and the μ TCB APIs. A full implementation would need to make this modification for a commercial product, but for feasibility-testing purposes it is sufficient to execute simple applications directly over *NWorld*.

System Boot. The μ TCB (running in *SWorld*) initializes prior to the smartphone’s OS (running in *NWorld*). The TrustZone platform allows μ TCB to specify which peripherals’ interrupts should be handled by the smartphone’s OS, and which by μ TCB. Initially, the only I/O devices handled by μ TCB are the SAK and the LED indicator, all other peripherals are handled in *NWorld* by the smartphone’s OS.

Switching to Secure-Mode and Interrupt Masking. When the user presses the SAK, μ TCB registers itself to handle interrupts from the touch-screen, which is required to interact with the user while the smartphone is in secure-mode (this is changed back when the user exits secure-mode). If the user requests to block a sensor for privacy reasons (e.g., microphone, camera), μ TCB masks its interrupts. The only full peripheral-device driver required to deploy in μ TCB is, therefore, that of the touch-screen.

Such a scheme efficiently handles interrupts in a TrustZone based implementation of μ TCB. However, its implementation is difficult: it requires careful bookkeeping of the touch-screen’s state which should be restored when the device exits secure-mode, and handling edge cases where the OS is currently using the peripherals that μ TCB needs to block. We leave addressing these issues in our prototype for future work, and avoid them in our evaluation below by using simple applications that execute directly over *NWorld*.

Architecture Overhead Evaluation. In order to evaluate our implementation, we utilized the ARM FastModels product that simulates an ARM device with the Cortex-A9 processor that is common in smartphones. We run the simulator on a PC with an Intel Core 2 Duo processor and 4GB of RAM, resulting in a wait time of roughly 80 seconds to load the simulation. In order to evaluate the overhead of μ TCB, we measure the number of additional processor cycles while performing different operations (system boot, API calls and interactions with the user) and use these measurements to estimate execution time on the Cortex-A9 processor.

The overhead introduced to the smartphone boot process is minimal: the μ TCB only initializes the interrupt handlers, and switches to NWorld to continue the normal boot sequence of the smartphone; the μ TCB initialization requires 1600 processor cycles (less than 3μ s on the Cortex-A9 processor).

The cost of switching between the normal and secure worlds is roughly 3000 cycles, or 5μ s on the Cortex-A9. However, it is only ‘paid’ when an application invokes the μ TCB via an API call or when the user presses the SAK to enter secure-mode, and not during the typical operation of the smartphone: When not in secure-mode, all I/O interrupts (except for the SAK) are handled in NWorld by the smartphone’s OS, and therefore do not require (expensive) switching between worlds.

Handling the SAK has the highest overhead, since the μ TCB ‘takes over’ the screen and displays the menu to the user, but requires less than 10^6 cycles or 2ms on the Cortex-A9, which is sufficiently short to remain unnoticeable to the user.

7 Conclusions

Cellular phones have evolved to ‘smartphones’, complex devices that are much more vulnerable to malware. This is especially alarming considering their widespread use and access to sensitive information, sensors data, video/audio feeds and the like.

In order to secure the complex smartphones, we have proposed to delegate critical security services from the smartphone’s OS to a small and separated μ TCB module. Our design integrates the established concepts of minimal TCB, trusted physical interfaces, user training and forcing ceremonies within the context of smartphones. We have demonstrated the use of μ TCB and SAK for two smartphone applications: messaging and mobile payments. Finally, we have presented a proof-of-concept (PoC) implementation of μ TCB by utilizing TrustZone, an existing mobile-phone trusted computing platform, and evaluated its performance using ARM’s device simulator.

During this work, we found that although modern OSs are made aware of underlying hardware security platforms, it may be challenging to use some of the platforms’ features in practice (e.g., TrustZone’s interrupt-handlers remapping). We hope that our work will motivate OS vendors to improve their support for such platforms in the face of increasing threats to users through their smartphones.

7.1 Future Work

Several significant steps are needed to bring our PoC implementation into a generally usable system. To deploy it with a fully-featured OS, one will need to patch the OS to allow efficient remapping of interrupt handlers and expose the μ TCB API library to applications. The simulator code must also be ported to smartphone hardware, a nontrivial engineering challenge. Finally, a usability study with deployment on fully-functional smartphones is important in order to evaluate and improve the proposed security indicator, user interfaces and the interaction protocols between the user and μ TCB.

This work also leaves several questions open for research. For one, it is unclear how to best present pending requests from applications to the user in secure-mode. It is also interesting to consider alternate secure-mode indicators, and effective feedback methods to train users to notice such an indicator. Finally, it would be interesting to leverage the smartphone’s sensors to derive keys without relying on text passwords.

Acknowledgments

We thank Idan Warszawski for his contribution to an early Arduino-based implementation of this work, and the anonymous referees for their thorough comments and suggestions.

References

- [1] A. Herzberg. Payments and Banking with Mobile Personal Devices. *CACM*, 46(5):53–58, 2003.
- [2] R. Weinmann. Baseband Attacks: Remote Exploitation of Memory Corruptions in Cellular Protocol Stacks. In *WOOT*, pages 12–21. USENIX Association, 2012.
- [3] Z. Xu and S. Zhu. Abusing Notification Services on Smartphones for Phishing and Spamming. In *WOOT*, pages 1–11. USENIX Association, 2012.
- [4] T. Vidas, D. Votipka, and N. Christin. All Your Droid Are Belong to Us: A Survey of Current Android Attacks. In *WOOT*, pages 81–90. USENIX Association, 2011.
- [5] M. Farb, Y. Lin, M. Burman, G. Chandok, J. McCune, and A. Perrig. SafeSlinger: Easy-to-Use and Secure Public-Key Exchange. <https://sparrow.ece.cmu.edu/group/pub/SafeSlinger.pdf>, 2012.
- [6] A. Vasudevan, E. Owusu, Z. Zhou, J. Newsome, and J. McCune. Trustworthy Execution on Mobile Devices: What Security Properties Can My Mobile Platform Give Me? In *TRUST*, pages 159–178. Springer, 2012.
- [7] V. Gligor, C. Chandrasekaran, R. Chapman, L. Dotterer, M. Hecht, W. Jiang, A. Johri, G. Luckenbaugh, and N. Vasudevan. Design and Implementation of Secure Xenix. *IEEE TSE*, pages 208–221, 1987.
- [8] J. McCune, A. Perrig, A. Seshadri, and L. Doorn. Turtles All The Way Down: Research Challenges in User-Based Attestation. In *HotSec*, 2007.
- [9] J. McCune, Y. Li, N. Qu, Z. Zhou, A. Datta, V. Gligor, and A. Perrig. TrustVisor: Efficient TCB Reduction and Attestation. In *IEEE Symposium on Security and Privacy*, pages 143–158. IEEE Computer Society, 2010.
- [10] N. Santos, H. Raj, S. Saroiu, and A. Wolman. Trusted Language Runtime (TLR): Enabling Trusted Applications on Smartphones. In *HotMobile*, pages 21–26. ACM, 2011.
- [11] Z. Zhou, V. Gligor, J. Newsome, and J. McCune. Building Verifiable Trusted Path on Commodity x86 Computers. In *IEEE Symposium on Security and Privacy*, pages 616–630. IEEE Computer Society, 2012.
- [12] H. Liu, S. Saroiu, A. Wolman, and H. Raj. Software Abstractions for Trusted Sensors. In *MobiSys*, pages 365–378. ACM, 2012.
- [13] H. Raj, S. Saroiu, A. Wolman, and J. Padhye. Splitting the Bill for Mobile Data with SIMlets. In *HotMobile*, pages 1–6. ACM, 2013.
- [14] C. Karlof, J. Tygar, and D. Wagner. Conditioned-Safe Ceremonies and a User Study of an Application to Web Authentication. In *NDSS*. The Internet Society, 2009.
- [15] A. Herzberg and R. Margulies. Forcing Johnny to Login Safely: Long-Term User Study of Forcing and Training Login Mechanisms. In *ESORICS*, pages 452–471. Springer, 2011.