

# A Study on Power Side Channels on Mobile Devices

Lin Yan, Yao Guo, Xiangqun Chen, Hong Mei

Key Laboratory of High-Confidence Software Technologies (Ministry of Education),  
School of Electronics Engineering and Computer Science, Peking University, Beijing, China  
{y.l, yaoguo, cherry, meih}@pku.edu.cn

## ABSTRACT

Power side channel is a very important category of side channels, which can be exploited to steal confidential information from a computing system by analyzing its power consumption. In this paper, we demonstrate the existence of various power side channels on popular mobile devices such as smartphones. Based on unprivileged power consumption traces, we present a list of real-world attacks that can be initiated to identify running apps, infer sensitive UIs, guess password lengths, and estimate geo-locations. These attack examples demonstrate that power consumption traces can be used as a practical side channel to gain various confidential information of mobile apps running on smartphones. Based on these power side channels, we discuss possible exploitations and present a general approach to exploit a power side channel on an Android smartphone, which demonstrates that power side channels pose imminent threats to the security and privacy of mobile users. We also discuss possible countermeasures to mitigate the threats of power side channels.

## Categories and Subject Descriptors

D.4.6 [Operating Systems]: Security and Protection—*Invasive software*

## General Terms

Measurement, Reliability

## Keywords

Mobile security, side channel attack, power side channel

## 1. INTRODUCTION

Side channel attacks have drawn a lot of research attention for a long time. Originally, side channel attacks refer to attacks aimed to break a cryptosystem relying on physical information, which is neither the plain-text nor

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*Internetware '15* November 6 2015, Wuhan, China  
Copyright 2015 ACM (TBD) ...\$15.00.

cipher-text. Recently, the context of side channel attacks has been expanded from cryptosystems to all kinds of computing systems. The goal of side channel attacks is gaining confidential information from the targeted computing system, while leveraging “side channel” information. Previously discovered side channels include timing information [5, 13], sound [57, 14], electro-magnetic radiation [44], shared memory/registers/files between processes [55, 53], sensor information [43, 41] and power consumption [24, 34], etc.

### 1.1 Power Side Channels

Power analysis attacks (or *power side channels*, PSCs) have become an important type of side channel attacks. The most famous example of power analysis attacks is the recovery of an encryption key from a cryptosystem [24, 23]. Messerges *et al.*[32] examined both *simple power analysis (SPA)* and *differential power analysis (DPA)* attacks against the data encryption standard algorithm and managed to breach the security of smart-cards using signal-to-noise ratio (SNR) based multi-bit attack.

On mobile devices such as smartphones, Michalevsky *et al.* proposed PowerSpy [34], which investigates the relationship between signal strength and the power pattern of smartphones. They showed that it is possible to infer smartphone users’ whereabouts based on the power traces.

Our work also focuses on the mobile platform, but we will present a more comprehensive study of different PSCs on mobile devices.

### 1.2 Attacks on Smartphones

Mobile devices including smartphones and tablets have become a popular computing platform. Mobile users are storing more and more privacy information such as passwords, credit card numbers, geo-locations, contacts information and even biometric information like fingerprints. Unfortunately, these sensitive data are vulnerable to various attacks, as evidenced by a huge number of malware aiming at stealing user privacy on Android devices [56].

One popular example is UI-based attacks. For example, ScreenMilker [29] can take screenshots of the foreground app covertly thus stealing user credentials. Chen *et al.* proposed an attack on the Android platform called UI inference attack [12]. They use the shared-memory side channel to infer UI states, in order to detect the correct timing for attacks.

There are many other attacks on mobile devices, we will present several of them that can be enabled by exploiting *only* power traces of a device.

### 1.3 Summary and Contributions

In this paper, we present our study on the existence of *power side channels* (PSCs) on mobile devices, which may bring serious threats to mobile user security and privacy. We also demonstrate some potential attacks based on these PSCs.

We conduct our study on PSCs through a series of experiments to collect power patterns of mobile devices. The power patterns can be collected through both hardware-based method (with a Monsoon Power Monitor [2] attached to the smartphone) and software-based method (directly polling voltage and current readings within the mobile system). Then we analyze collected power traces to see if the uniqueness of different characteristics of the confidential data can be reflected on power patterns.

With experiments on a Nexus 5 smartphone, we demonstrate several prominent PSCs on mobile devices that can be exploited to gain various kinds of confidential information. Besides previously reported PSCs used to estimate the geo-location of mobile devices [34], we also identify the following new PSCs that can be used to: (1) identify different mobile apps (based on power patterns during app loading); (2) infer different UIs within one mobile app (based on collected sensitive UI traces); and (3) guess password lengths (based on power patterns during password input).

Based on these newly discovered PSCs, we discuss possible exploitations and present a comprehensive exploitation scenario, in which the malicious party could steal user passwords by leveraging PSCs in a particular order. To show how the exploitation works in reality, we present an overview of a general approach to exploit PSCs on Android smartphones.

Our contributions can be summarized as follows.

- We present the first demonstrative study to reveal the existence of multiple power side channels on smartphones.
- We demonstrate that power side channels pose practical threats as they can be used to fulfill different attack purposes, such as app identification, UI identification or password guessing.
- We present a general approach to exploit power side channels.
- We discuss possible countermeasure techniques to mitigate the threats of power side channels.

The rest of this paper is organized as follows. We first introduce our methodology in Section 2. In section 3, we present various power side channels we discover on mobile devices. Section 4 presents the details of our research on the exploitation of power side channels. We discuss our limitations and possible defenses and future work in Section 5. Background of our work and survey of related work are presented in Section 6 and Section 7 concludes this paper.

## 2. METHODOLOGY

In this section, we introduce the experimental setup and how we collect power traces from a smartphone.

### 2.1 Experimental Setup

We use a Google Nexus 5 smartphone to perform most of the experiments. The power numbers can be measured with

Table 1: Battery status files used in our experiments (on the Android platform for Nexus 5).

File location	Description	Unit
/sys/class/power_supply/battery/voltage_now	Instant voltage reading of battery	$\mu V$
/sys/class/power_supply/battery/current_now	Instant current reading of battery	$\mu A$

a Monsoon Power Monitor with a stable 4.2V voltage. We install popular apps on Android ASOP 4.4 on the testing phone and measure the power numbers in different cases.

### 2.2 Hardware Measurement

Hardware measurement is straightforward. We use the Monsoon Power Monitor [2] in our study to supply a stable voltage to smartphones and measure the current value in various stages. The power monitor works as an accurate hardware power meter, which offers stabilized power supply to the smartphone being tested and records the power consumption numbers of the smartphone in real-time. Since the voltage is a constant, we simply use current value to denote the power consumption.

### 2.3 Software Measurement

Although power measurements based on hardware meters can be used to confirm the existence of PSCs, it might still not be applicable to real-world exploitation since the hardware meter is apparently not a standard auxiliary accessory for smartphones.

Fortunately, power related statistics are publicly accessible on most smartphone OSes such as Android. In general, instant power numbers can be calculated based on voltage and current readings of BMU (Battery Monitoring Unit) [51]. The battery status information is accessible by most apps without system-level privilege.

In our experiments, we develop a polling collector on a Nexus 5 smartphone which reads and records power patterns by polling battery status files, which are listed in Table 1.

### 2.4 Identification of Power Side Channels

In this paper, a *power side channel* (PSC) is defined as an identifiable power pattern that can be used to infer the secret information from a running app on a mobile device.

We assume that power consumption traces are available as an unprivileged source that can be obtained by any normal apps running on a mobile device, which is the case for most Android devices.

The purpose of this paper is to demonstrate the existence of PSCs on mobile devices that can be exploited to perform various attacks. We achieve this by providing experimental results that show repeatable and unique power traces that can be detected during runtime. We also cover the details on how to exploit these PSCs to conduct real attacks. We believe that these PSCs pose real threats to the security and privacy of mobile users.

## 3. POWER SIDE CHANNELS

This section presents our findings on potential power side channels (PSCs) on mobile devices, especially three new

PSCs that can be exploited to identify running apps, infer sensitive UIs and guess password information, respectively.

### 3.1 PSC #1: App Identification

For many real-world attacks aiming at mobile systems, knowing which app is running in the foreground is important. Knowing the identity of the foreground app means being aware of the timing information of the attacking target. Previously, identity of the foreground app can be accessed by calling the `getRunningTasks()` method of the `ActivityManager` on Android for all apps. However, this method is officially unavailable to third-party apps after Android 5.0 due to security concerns [1].

Based on our findings, the foreground app can also be identified by analyzing the power patterns of the app loading phase. We choose three popular apps from three common app categories (social media, e-commerce and productivity) and record the power patterns of the loading phases. The results are shown in Figure 1. It is obvious that different apps exhibit distinguishable power patterns and different test runs of the same app generate visually similar power curves.

We have conducted more than two dozen different test runs on these apps to test the repeatability of the power patterns during app loading. The results show similar patterns as the data shown in Figure 1. However, due to space limit, we cannot show all the results in this paper.

On the other hand, we have also tested the app to collect its power pattern when the apps continue to run after the app loading phase. The results show that the power patterns during the app loading phase is distinguishable from power traces in other execution phases.

We also collected power traces during app loading using the software-based method and got similar results. Because software-based power traces are more practical for real-world attacks, we will present only software-based traces in the next case.

### 3.2 PSC #2: UI Inference

The UI states within a mobile app are also crucial for many attacks. For example, in order to initiate activity hijacking attacks on Android, attackers need to know when the user login UI will be prompted so that they can intercept the UI state transitions and insert fake user login UIs that could steal user credentials.

However, what is going on inside an app is not directly observable by other parties without system-level privileges if the app chooses not to share it. A previous study has shown that attackers could peek into mobile applications via the shared-memory side channel along with some UI signatures based on CPU utilization time and other runtime events [12].

Different UI states within one mobile app can also be inferred through their respective power traces. As depicted in Figure 2, we conduct experiments on collecting power patterns of visiting three different UIs in the Amazon app in multiple test runs. The power patterns demonstrate both the repeatability of a single UI and the distinguishability between different UIs within the same mobile app.

For example, in order to infer the login screen of a mobile app, an attacker can first collect the power patterns of the targeted UI in advance on another device. After learning the power pattern, it is possible to detect the pattern on the targeted device with a pattern matching algorithm. In this

way, the attacker would be able to infer when the targeted UI is loaded on the targeted device, i.e., recognizing the exact timing to conduct attack.

### 3.3 PSC #3: Password Length Guessing

For a more complex scenario, we attempt to explore whether power traces can be used to reveal information on user passwords, similar to previous power analysis approaches on other platforms.

Although it is difficult to guess the actual passwords, we find that it is relatively easy to identify the length of a password during password input. Figure 3 depicts the hardware-measured power traces during password input. Each time the user inputs a single letter of the passwords, it will be reflected in the power pattern as a sudden “burst” of power consumption. It is obvious that the length of user’s password can be inferred by counting the number of consecutive power pattern “bursts”.

Obtaining the length of a user password is meaningful for attackers to guess the password since it reduces the complexity of the cracking algorithm, as well as the size of the helping dictionary.

### 3.4 PSC #4: Geo-location Estimation

The geo-location of a mobile device can also be revealed by its power consumption pattern. A recent study uses machine learning techniques to identify the routes taken by the smartphone users based on previously collected power consumption data [34].

There are many directions can be followed on this thread of work. For example, based on our observation, the power consumption patterns can be used to track the accurate real-time location of a mobile device, if a database of power patterns on location signatures has been collected in advance.

## 4. PSC EXPLOITATIONS

In this section we present a detailed description of how to exploit power side channels (PSCs) on mobile devices. We first discuss possible exploitations on mobile devices leveraging presented PSCs. Then we introduce a general PSC Exploitation approach.

### 4.1 Possible PSC Exploitations

PSCs on mobile devices can be exploited by malicious parties to steal user privacy information such as login credentials and geo-locations.

For example, as depicted in Figure 4, PSCs #1, #2 and #3 found in the previous section can be combined to guess user login credentials of a certain app. To be more specific, the attack can first learn the power patterns of the victim app, the target UI and the user input. Then the attacker can develop a malicious app to fulfill PSC exploitations. The malicious app can first detect the loading of the victim app via PSC #1; then the malicious app can capture the occurrence of the login UI via PSC #2; finally the malicious app can help crack user password with the password length information acquired via PSC #3. It should be noted that the malicious app does not need any special permission to read system power patterns on popular mobile systems like Android.

The key of the malicious apps to exploit PSCs is to achieve automated detection of some pre-learned power patterns. This can be done by adopting pattern matching or machine

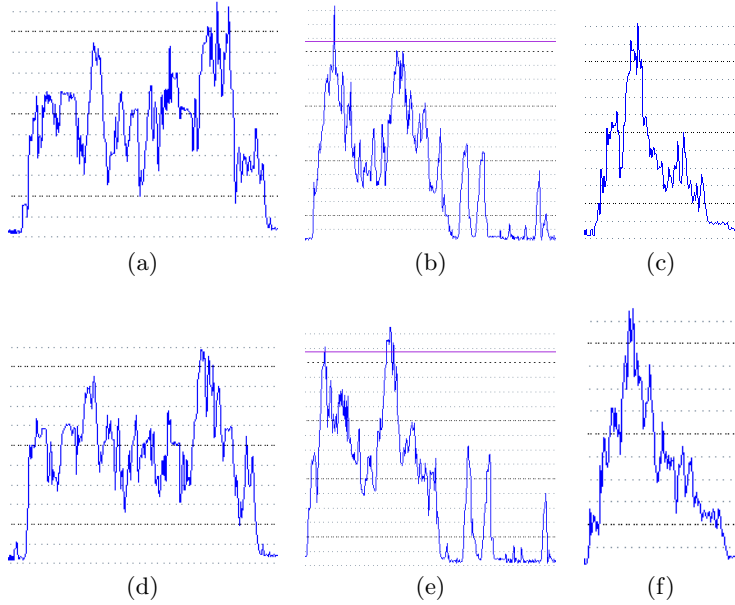


Figure 1: Hardware-based measurements on the loading phases of different apps in two test runs: (a), (d) for WeChat; (b), (e) for Alipay; (c), (f) for Gmail.

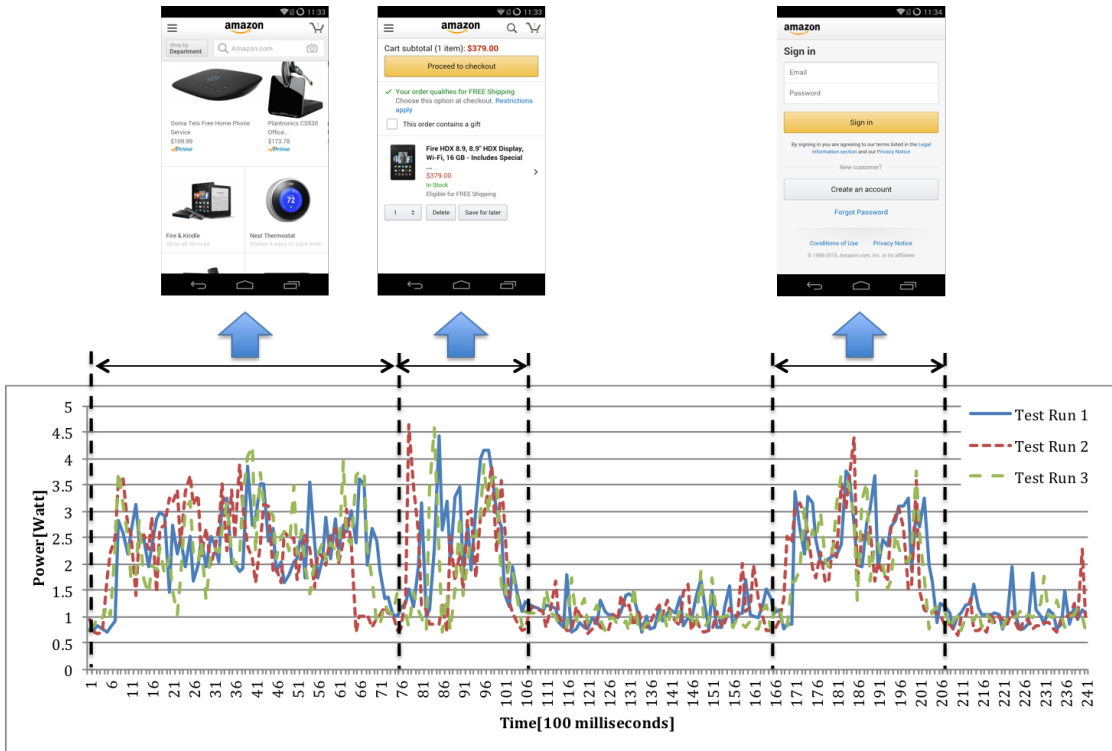


Figure 2: Power traces collected through software-based measurement for the Amazon app. The three UIs shown above are *app entry*, *product details*, and *user login*, respectively.

learning algorithms, such as dynamic time warping (DTW) [37], which can be used to match a previously learned pattern in a continuous power trace. However, the pattern matching process may affect the total system power consumption during detection, thus it should be implemented

as a light-weight detection process if used in real-time.

In the above exploitation scenario, all attacks are based on PSCs. However, it is not necessary to use the power channel alone to conduct the attack. Attackers can choose to combine PSCs with other attacking channels to improve the

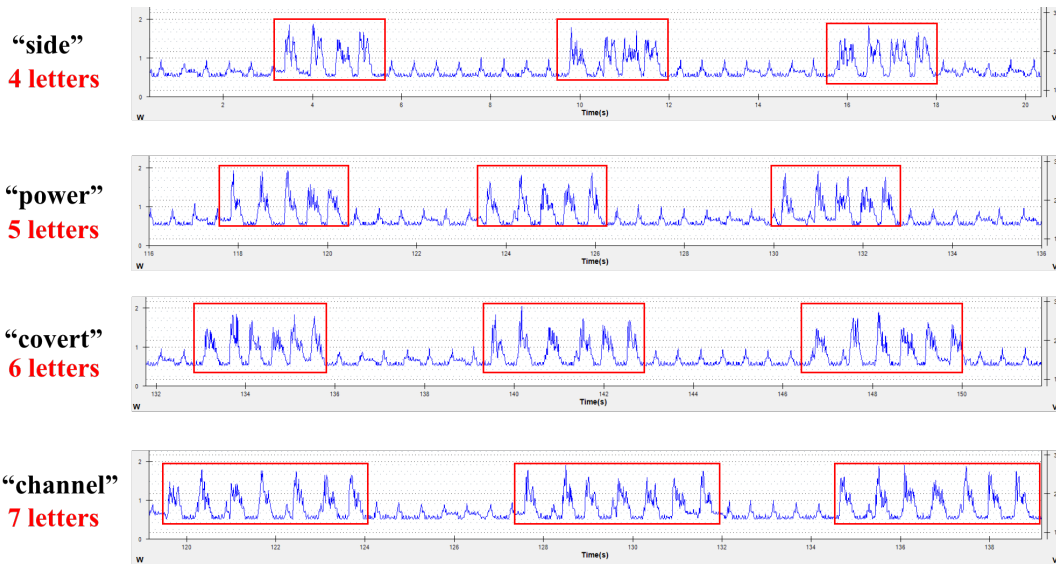


Figure 3: Power patterns from hardware measurements when inputting passwords with different lengths.

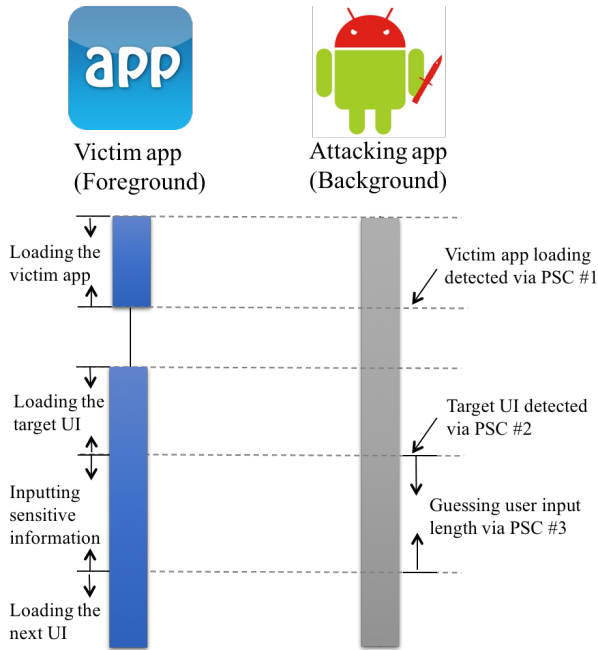


Figure 4: Possible Exploitations of presented PSCs.

attacking success rates. On particular example is that an attacker can choose to take screenshots continuously after detecting the target UI and apply image recognition algorithms to retrieve user passwords in the screenshots. Apparently, screenshot-based attacks are orthogonal to PSC based attacks we discussed and out of scope of this paper.

## 4.2 A General Exploitation Approach

Figure 5 shows the overview of a general PSC exploitation approach. The purpose of PSC exploitation is to find an effective method to detect the occurrences of the target power pattern from a continuous power trace collected from an app. PSC exploitation involves the following main steps.

### 4.2.1 Training Data Collection

The first step in collecting training data is selecting a target. The choice of the target depends on which kind of PSC the attacker wants to exploit. For example, the loading phase of a mobile app might be the target if the attacker wants to leverage PSC # 1 described in Section 3.1. Or if the attacker wants to leverage PSC # 2 as shown in Section 3.2, the loading of sensitive UIs are considered to be the target.

After choosing the target, the attacker must specify the start and end of the target explicitly so that its power pattern can be precisely measured. As we have introduced in Section 2.3, the voltage and current numbers of running the target can be acquired by polling unprivileged system power files and the power can be calculated based on them.

### 4.2.2 Model Training

Based on the collected power traces of the target, the attacker needs to generate a signature to represent the target. This step generally involves detecting features of the collected power traces and signatures accordingly. The specific methods for detecting features abstracting the signature depend on the characteristics of the training dataset and PSC exploitation scenario.

### 4.2.3 Target Detection

We assume that an attacking app runs in the background on the smartphone, which continuously collect the power traces. Based on the signature of the target trained in previous steps, the attacking app is able to detect the occurrence of the target in real-time by running appropriate detection algorithms. Detection algorithms are pattern matching algorithms in essence and the choice of them depends on characteristics of the PSC exploitation scenario.

## 5. DISCUSSIONS

This section discusses the limitations and countermeasures on the PSCs identified in the previous section.

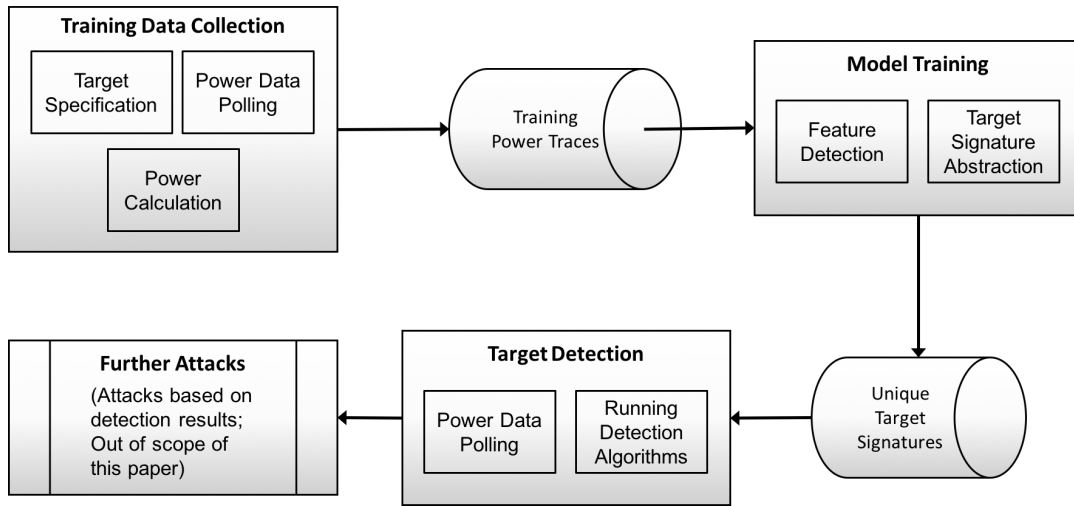


Figure 5: An overview of a general PSC exploitation approach.

## 5.1 Limitations

Although we have demonstrated the existence of multiple PSCs. There are some limitations in the current study, most of which we plan to explore in future work:

- Identifying the side channel is only the first step. In order to perform real-world attacks using these side channels, there are various issues that should be considered, such as how to collect power patterns accurately, how to detect app loading or sensitive UIs in real-time, how to apply successful attacks in practical scenarios, etc.
- One particular issue that needs to investigate is how to guarantee that the power patterns keep unchanged during real-time detection. For example, a pattern matching app should be carefully implemented such that it does not cause significant power overheads during real-time detection. Otherwise it will difficult to identify the power patterns due to power overhead of the detecting process itself.
- In our experiments, we have only tested the ideal case in which the background power consumption noise is minimal. In reality, there might exist multiple apps running in the background. Because these background apps might make network or I/O requests in an irregular way, it may affect the power patterns when we tried to detect the patterns revealed by the PSCs. This issue may also increase the difficulty to exploit discovered PSCs, nonetheless PSCs are practical threats on mobile devices.

## 5.2 Possible Mitigating Techniques

- **Energy obfuscation through code injection.** One straightforward mitigation approach is that we can inject meaningless code into mobile apps, in order to insert power bursts into its power pattern to make it unpredictable. This can be achieved at the source-code level during app development, or through instrumentation to the bytecode for app binaries.

- **Randomly changing display/color parameters.** One interesting feature for the OLED or AMOLED displays used for smartphones is that it consumes different power when different color schemes are used [15]. Thus we can vary the displaying color and other parameters during the execution of apps we want to protect. This could also be achieved during app development or through bytecode instrumentation [28].
- **Raising the privilege needed to access power files.** Of course, we can always make the power information privileged, such that not all apps could access these data directly. As a matter of fact, mobile apps probably do not need to read low-level power related files containing raw voltage or current readings. The only thing that most apps need to know is how much battery is still remaining, which should not pose serious threats as a side channel.

## 6. BACKGROUND AND RELATED WORK

In this section, we introduce background knowledge on power characteristics of Android smartphones and related research work on side channel attacks, especially power side channel attacks.

### 6.1 Power Characteristics of Android Smartphones

In this paper, we use Android smartphones as our studying subject as Android is currently the most popular mobile operating system for smartphones and tablets. Similar algorithms could be applied on Android tablets, as they share the same framework and most of the APIs.

For a different mobile OS such as Apple iOS, the power consumption patterns of their mobile apps share very similar characteristics as Android because they use similar hardware components and similar app structure. Thus the attacks could potentially be repeated on iPhones as well.

Most smartphones are power-hungry devices, while many batteries could last only less than a day in typical usage. The most power consuming components are CPU, network, screen and various sensors such as GPS and camera [10]. For

most Android smartphones, the screen consumes a majority of the total power. The screen power could be affected significantly due to its brightness and color of pixels [15], thus it is able to reduce the power consumption of an app by adjusting its color schemes [27].

For a mobile application, the power of a particular execution trace in a certain stable environment is determined by the behavior of the app. Different types of apps consume different power according to their usage of CPU cycles, network traffic and screen brightness, etc. Power consumption for an app can therefore be modeled by their resource usages [51], system call traces [50] and source code [17], etc.

Besides power consumed when an app runs in the foreground (i.e., with user interaction), many apps also consume significant power when their respective services run in the background, for example, checking new emails or new tweets regularly. Although background power could potential drain a lot of battery in the long run, it is relative small compared to the power consumed by foreground apps. In this paper, we only consider the power patterns of an app when it runs in the foreground.

## 6.2 Side Channel Attacks

Side channel attacks have drawn researchers' attention for decades. Wright and Greengrass [48] reported the first official record on utilizing side channel attacks that happened in 1965. At first, research on side channel attacks mainly focused on cryptography [25, 4, 23] and encrypted communications [42, 8, 40, 47]. In these study, side channel attacks are defined as successfully decrypting the ciphertext based on the information gathered from the encryption system that is neither the plain-text being encrypted nor the cipher-text after encryption.

Along with the research development, side channel attacks are discussed more broadly in many contexts, not necessarily limited to the contexts of cryptography and encrypted communications. Previously discovered common side channels include timing channels [46, 16, 9, 5, 13, 26, 54], acoustic channels [57, 14], electromagnetic waves [3, 44], shared memory/registers/files between processes [38, 20, 55, 53], sensor metrics [49, 35, 43, 41] and power consumption [24, 34]. Types of information can be leaked via these side channels include almost anything related to sensitive information about users, such as keystrokes [21], locations [34], speech [33] or even health data [13].

In order to protect computing systems from side-channel attacks, many research work have proposed various mitigation or protection methods. Examples include redesigning the encryption methods to prevent power analysis [36, 7], predictive mitigation of timing channels for interactive applications [52], system-level protection against cache-based side channel attacks [22], and new cache designs to thwart software cache-based side channel attacks [45].

## 6.3 Power Side Channels

Power analysis attacks (or power side channels) [6] have become an important type of side channel attacks in recent years. One well-known example of power analysis is the recovery of an encryption key from a cryptosystem [24, 23]. Messerges *et al.* [32, 31] examined both simple power analysis (SPA) and differential power analysis (DPA) attacks against the data encryption standard (DES) algorithm and managed to breach the security of smart-cards using the pro-

posed signal-to-noise ratio (SNR) based multi-bit attack.[30]

In order to defeat power analysis attacks on smart-cards, Herbst *et al.* [18] present an efficient AES software implementation that is suited for 8-bit smart-cards and resistant against power analysis attacks. Ratanpal *et al.* [39] presents a circuit that can be added to crypto-hardware to suppress information leakage through the power supply pin side channel.

Chari *et al.* [11] propose a sound approach to counter-act power analysis attacks. It includes an abstract model which approximates power consumption in most devices and a generic technique to create provably resistant implementations for devices where the power model has reasonable properties. They prove a lower bound on the number of experiments required to mount statistical attacks on devices whose physical characteristics satisfy reasonable properties.

Power side channels have also been discovered on other systems besides smart-cards. For example, Hlavacs *et al.* [19] demonstrate that energy consumption side-channel attack can be performed between virtual machines in a cloud.

On mobile platforms, Michalevsky *et al.* proposed PowerSpy [34], which investigates the relation between signal strength and the power pattern of the smartphone and showed that they can infer smartphone users' whereabouts based on the power traces.

Our work also focuses on the mobile platform, but we have presented different and more general exploitations based on power traces.

## 7. CONCLUDING REMARKS

In this paper, we have demonstrated the existence of various power side channels that can be exploited to perform attacks to steal private information about apps running on mobile devices.

We first provide our discovery of several power side channels on mobile devices. Then we provide a general exploitation approach of these power side channels.

We believe these power side channels pose real threats to mobile user security and privacy if they are exploited in real-world attacks, thus needed to be addressed in an urgent manner.

## 8. ACKNOWLEDGEMENT

This work is supported by the High-Tech Research and Development Program of China under Grant No. 2013AA01A605 and the National Natural Science Foundation of China under Grant No. 61421091 and No. 61103026.

## 9. REFERENCES

- [1] Android developers - ActivityManager refernces. [http://developer.android.com/reference/android/app/ActivityManager.html#getRunningTasks\(int\)](http://developer.android.com/reference/android/app/ActivityManager.html#getRunningTasks(int)).
- [2] Monsoon Power Monitor. <https://www.msoon.com/LabEquipment/PowerMonitor/>.
- [3] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi. The EM side channel (s). In *Cryptographic Hardware and Embedded Systems-CHES 2002*, pages 29–45. Springer, 2003.
- [4] E. Biham and A. Shamir. Differential fault analysis of secret key cryptosystems. In *Advances in Cryptology (CRYPTO '97)*, pages 513–525. Springer, 1997.

- [5] A. Bortz and D. Boneh. Exposing private information by timing web applications. In *Proceedings of the 16th International Conference on World Wide Web*, pages 621–628. ACM, 2007.
- [6] E. Brier, C. Clavier, and F. Olivier. Correlation power analysis with a leakage model. In *Cryptographic Hardware and Embedded Systems-CHES 2004*, pages 16–29. Springer, 2004.
- [7] E. Brier and M. Joye. Weierstraß elliptic curves and side-channel attacks. In *Public Key Cryptography*, pages 335–345. Springer Berlin Heidelberg, 2002.
- [8] D. Brumley and D. Boneh. Remote timing attacks are practical. In *Proceedings of the 12th Conference on USENIX Security Symposium - Volume 12*. USENIX Association, 2003.
- [9] S. Cabuk, C. E. Brodley, and C. Shields. Ip covert timing channels: design and detection. In *Proceedings of the 11th ACM conference on Computer and communications security*, pages 178–187. ACM, 2004.
- [10] A. Carroll and G. Heiser. An analysis of power consumption in a smartphone. In *USENIX annual technical conference*, volume 14, 2010.
- [11] S. Chari, C. S. Jutla, J. R. Rao, and P. Rohatgi. Towards sound approaches to counteract power-analysis attacks. In *Advances in Cryptology (CRYPTO 99)*, pages 398–412. Springer, 1999.
- [12] Q. A. Chen, Z. Qian, and Z. M. Mao. Peeking into your app without actually seeing it: UI state inference and novel Android attacks. In *Proceedings of the 23rd USENIX Conference on Security Symposium*, pages 1037–1052, Berkeley, CA, USA, 2014. USENIX Association.
- [13] S. Chen, R. Wang, X. Wang, and K. Zhang. Side-channel leaks in web applications: A reality today, a challenge tomorrow. In *Proceedings of the 2010 IEEE Symposium on Security and Privacy*, pages 191–206. IEEE Computer Society, 2010.
- [14] A. Das, N. Borisov, and M. Caesar. Do you hear what I hear?: Fingerprinting smart devices through embedded acoustic components. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, CCS '14*, pages 441–452. ACM, 2014.
- [15] M. Dong and L. Zhong. Power modeling and optimization for OLED displays. *IEEE Transactions on Mobile Computing*, 11(9):1587–1599, Sept 2012.
- [16] S. Gianvecchio and H. Wang. Detecting covert timing channels: an entropy-based approach. In *Proceedings of the 14th ACM conference on Computer and communications security*, pages 307–316. ACM, 2007.
- [17] S. Hao, D. Li, W. G. Halfond, and R. Govindan. Estimating mobile application energy consumption using program analysis. In *35th International Conference on Software Engineering (ICSE)*, pages 92–101. IEEE, 2013.
- [18] C. Herbst, E. Oswald, and S. Mangard. An AES smart card implementation resistant to power analysis attacks. In *Applied cryptography and Network security*, pages 239–252. Springer, 2006.
- [19] H. Hlavacs, T. Treutner, J.-P. Gelas, L. Lefevre, and A.-C. Orgerie. Energy consumption side-channel attack at virtual machines in a cloud. In *Dependable, Autonomic and Secure Computing (DASC), 2011 IEEE Ninth International Conference on*, pages 605–612. IEEE, 2011.
- [20] S. Jana and V. Shmatikov. Memento: Learning secrets from process footprints. In *2012 IEEE Symposium on Security and Privacy (S&P '12)*, pages 143–157, May 2012.
- [21] K. Killourhy and R. Maxion. Comparing anomaly-detection algorithms for keystroke dynamics. In *IEEE/IFIP International Conference on Dependable Systems Networks, 2009(DSN '09)*, pages 125–134, June 2009.
- [22] T. Kim, M. Peinado, and G. Mainar-Ruiz. Stealthmem: System-level protection against cache-based side channel attacks in the cloud. In *USENIX Security symposium*, pages 189–204, 2012.
- [23] P. Kocher, J. Jaffe, and B. Jun. Introduction to differential power analysis and related attacks. <http://www.cryptography.com/resources/whitepapers/DPATechInfo.pdf>, 1998.
- [24] P. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *Advances in Cryptology (CRYPTO '99)*, pages 388–397. Springer, 1999.
- [25] P. C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *Advances in Cryptology (CRYPTO '96)*, pages 104–113. Springer, 1996.
- [26] B. Köpf and G. Smith. Vulnerability bounds and leakage resilience of blinded cryptography under timing attacks. In *Proceedings of the 2010 23rd IEEE Computer Security Foundations Symposium, CSF '10*, pages 44–56. IEEE Computer Society, 2010.
- [27] D. Li, A. H. Tran, and W. G. Halfond. Making web applications more energy efficient for OLED smartphones. In *Proceedings of the 36th International Conference on Software Engineering*, pages 527–538. ACM, 2014.
- [28] D. Li, A. H. Tran, and W. G. J. Halfond. Making web applications more energy efficient for OLED smartphones. In *ICSE 2014*, pages 527–538, 2014.
- [29] C.-C. Lin, H. Li, X. Zhou, and X. Wang. Screenmilk: How to milk your Android screen for secrets. In *Proceedings of The 21th Annual Network and Distributed System Security Symposium (NDSS)*, 2014.
- [30] S. Mangard, E. Oswald, and T. Popp. *Power analysis attacks: Revealing the secrets of smart cards*, volume 31. Springer Science & Business Media, 2008.
- [31] T. S. Messerges, E. A. Dabbish, and R. H. Sloan. Power analysis attacks of modular exponentiation in smartcards. In *Cryptographic Hardware and Embedded Systems*, pages 144–157. Springer, 1999.
- [32] T. S. Messerges, E. A. Dabbish, and R. H. Sloan. Examining smart-card security under the threat of power analysis attacks. *IEEE Transactions on Computers*, 51(5):541–552, 2002.
- [33] Y. Michalevsky, D. Boneh, and G. Nakibly. Gyrophone: Recognizing speech from gyroscope signals. In *Proceedings of 23rd USENIX Security Symposium, USENIX Association*, 2014.
- [34] Y. Michalevsky, G. Nakibly, A. Schulman, and D. Boneh. PowerSpy: Location tracking using mobile device power analysis. In *24th USENIX Security*



- Symposium (USENIX Security 15)*, Washington, D.C., Aug. 2015. USENIX Association.
- [35] E. Miluzzo, A. Varshavsky, S. Balakrishnan, and R. R. Choudhury. Tappprints: Your finger taps have fingerprints. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services*, MobiSys '12, pages 323–336. ACM, 2012.
- [36] B. Möller. Securing elliptic curve point multiplication against side-channel attacks. In *Information Security*, pages 324–334. Springer, 2001.
- [37] M. Müller. Dynamic time warping. *Information retrieval for music and motion*, pages 69–84, 2007.
- [38] Z. Qian, Z. M. Mao, and Y. Xie. Collaborative TCP sequence number inference attack: How to crack sequence number under a second. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12, pages 593–604. ACM, 2012.
- [39] G. B. Ratanpal, R. D. Williams, and T. N. Blalock. An on-chip signal suppression countermeasure to power analysis attacks. *Dependable and Secure Computing, IEEE Transactions on*, 1(3):179–189, 2004.
- [40] T. S. Saponas, J. Lester, C. Hartung, S. Agarwal, and T. Kohno. Devices that tell on you: Privacy trends in consumer ubiquitous computing. In *Proceedings of 16th USENIX Security Symposium on USENIX Security Symposium*, pages 5:1–5:16. USENIX Association, 2007.
- [41] R. Schlegel, K. Zhang, X.-y. Zhou, M. Intwala, A. Kapadia, and X. Wang. Soundcomber: A stealthy and context-aware sound trojan for smartphones. In *Proceedings of The 18th Annual Network and Distributed System Security Symposium (NDSS)*, volume 11, pages 17–33, 2011.
- [42] D. X. Song, D. Wagner, and X. Tian. Timing analysis of keystrokes and timing attacks on SSH. In *Proceedings of the 10th Conference on USENIX Security Symposium - Volume 10*. USENIX Association, 2001.
- [43] R. Templeman, Z. Rahman, D. Crandall, and A. Kapadia. PlaceRaider: Virtual theft in physical spaces with smartphones. In *Proceedings of The 20th Annual Network and Distributed System Security Symposium (NDSS)*, Feb. 2013.
- [44] M. Vuagnoux and S. Pasini. Compromising electromagnetic emanations of wired and wireless keyboards. In *Proceedings of the 18th Conference on USENIX Security Symposium*, SSYM'09, pages 1–16. USENIX Association, 2009.
- [45] Z. Wang and R. B. Lee. New cache designs for thwarting software cache-based side channel attacks. In *ACM SIGARCH Computer Architecture News*, volume 35, pages 494–505. ACM, 2007.
- [46] J. C. Wray. An analysis of covert timing channels. In *Research in Security and Privacy, 1991. Proceedings., 1991 IEEE Computer Society Symposium on*, pages 2–7. IEEE, 1991.
- [47] C. V. Wright, L. Ballard, S. E. Coull, F. Monrose, and G. M. Masson. Spot me if you can: Uncovering spoken phrases in encrypted VoIP conversations. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, pages 35–49. IEEE Computer Society, 2008.
- [48] P. Wright and P. Greengrass. *Spycatcher: The candid autobiography of a senior intelligence officer*. Dell Publishing Company, 1987.
- [49] Z. Xu, K. Bai, and S. Zhu. TapLogger: Inferring user inputs on smartphone touchscreens using on-board motion sensors. In *Proceedings of the Fifth ACM Conference on Security and Privacy in Wireless and Mobile Networks*, WISEC '12, pages 113–124. ACM, 2012.
- [50] Y. Yetim, S. Malik, and M. Martonosi. EPROF: An energy/performance/reliability optimization framework for streaming applications. In *Design Automation Conference (ASP-DAC), 2012 17th Asia and South Pacific*, pages 769–774. IEEE, 2012.
- [51] C. Yoon, D. Kim, W. Jung, C. Kang, and H. Cha. AppScope: Application energy metering framework for Android smartphone using kernel activity monitoring. In *USENIX Annual Technical Conference*, pages 387–400, 2012.
- [52] D. Zhang, A. Askarov, and A. C. Myers. Predictive mitigation of timing channels in interactive systems. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 563–574. ACM, 2011.
- [53] K. Zhang and X. Wang. Peeping tom in the neighborhood: Keystroke eavesdropping on multi-user systems. In *Proceedings of the 18th Conference on USENIX Security Symposium*, SSYM'09, pages 17–32. USENIX Association, 2009.
- [54] Y. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart. Cross-VM side channels and their use to extract private keys. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12, pages 305–316. ACM, 2012.
- [55] X. Zhou, S. Demetriou, D. He, M. Naveed, X. Pan, X. Wang, C. A. Gunter, and K. Nahrstedt. Identity, location, disease and more: Inferring your secrets from Android public resources. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer Communications Security*, CCS '13, pages 1017–1028. ACM, 2013.
- [56] Y. Zhou and X. Jiang. Dissecting android malware: Characterization and evolution. In *IEEE S&P*, pages 95–109. IEEE, 2012.
- [57] L. Zhuang, F. Zhou, and J. D. Tygar. Keyboard acoustic emanations revisited. *ACM Transaction Information System Security*, 13(1):3:1–3:26, Nov. 2009.