MagAttack: Remote App Sensing with Your Phone

Zhuangdi Zhu
Nanjing University
zhuangdizhu@yahoo.com

Xiaoyu Ji
HKUST
xji@cse.ust.hk

Hao Pan
Shanghai Jiao Tong University
phoena_hp@yahoo.com

Fan Zhang
HKUST
fzhangee@connect.ust.hk

Yi-Chao Chen
UT Austin
yichao@cs.utexas.edu

Chuang-Wen You
National Taiwan University
cwyou@ntu.edu.tw

Figure 1: MagAttack Implementation Prototype.

UbiComp/ISWC '16 Adjunct, September 12-16, 2016, Heidelberg, Germany
ACM 978-1-4503-4462-3/16/09.
http://dx.doi.org/10.1145/2968219.2971404

Abstract
By tracking changes in electromagnetic radiation footprint emitted from computers using a magnetometer on commodity mobile devices, a malicious attacker can easily learn the secrets of the computer's owner without physically peeping at or hacking into the victim's system. Targeting at Applications and Web browsers, we present MagAttack, which uses the built-in magnetometer on commodity mobile phones to infer which App is running or which webpage the user is browsing on a nearby computer, as well as finer-grained information about victim's interests, habits, etc.

Our preliminary results show that MagAttack is independent of the earth's magnetic filed, model of phones, and magnetometer sampling rates. We also conducted an in-the-wild evaluation where an instrumented participant uses her laptop as usual and MagAttack can detect when she opens 10 different popular Apps. MagAttack achieves a classification accuracy of up to 98%.

Author Keywords
Side Channel Attack; Magnetometer; Commodity Mobile Device.

ACM Classification Keywords
D.4.6 [OPERATING SYSTEMS]: Security and Protection

Introduction

Modern mobile devices have many sensors that enable rich user experience. Being generally put to good use, they can be sometimes used by malicious attackers as handy tools to retrieve nearby target's private information. While the privacy risks associated with sensors like microphones, cameras, and WiFi receivers are obvious and well understood, the other risks remained under the radar for users. In this poster, we present MagAttack, an app on Android or iOS phones that can capture electromagnetic radiation footprint emitted from nearby computers to infer private information.

When an App starts to run on a computer, a set of instructions are executed by the computer's CPU and the electric current through the CPU circuit generates unique electromagnetic radiations (EMRs). Due to the low sampling rate (up to 200Hz) on most mobile devices, it's difficult to capture the details of the running CPU instructions. However, we show that, by extracting both time and frequency domain features from EMRs, MagAttack can still learn the footprint with a low sampling rate and can infer which App the victim is interacting with. The learned footprint is independent of the earth's magnetic filed, the model of phones, or the magnetometer sampling rate. MagAttack can distinguish 10 popular Apps that are available in Mac OS and achieve up to 98% classification accuracy.

Related Work

Side Channel Attack is to extract sensitive information surreptitiously from a system based on the system's physical implementation. Attacks can be conducted from various channels, including timing analysis [3], power consumption [2], electromagnetic radiation [7] and device motion [4].

The electromagnetic field is one of the important channels for information leaks. In [2], the authors utilize a single-point monitor to classify the usage of electrical alliances in home. They find the electromagnetic radiations generated by switch mode power supplies (SMPS) to be appliance-dependent. Another work [6] utilizes wearable devices to detect the electromagnetic strength around for the health concern. Some other works [7, 1] utilize high-frequency sensors to track the EMRs generated by different CPU operations. They focus on the decryption of popular cryptographic implementations in CPU, such as RSA and AES.

We differ from these works by that we are the first to use the magnetometer on commodity mobile devices to infer running apps or browsed web sites on nearby computers.

Methodology

We find that when opening or closing an application, the laptop CPU emits EMRs that distinct to each application. In order to collect the electromagnetic traces to analyze which application the laptop is running, we put a mobile phone near the CPU of the target laptop when it is running a background script to open and close different applications iteratively. In the mobile phone, we use its inside 3-axis magnetometer to track the surrounding EMRs. The EMRs are measured in in unit of microtesla (μT) along each axis and then square rooted.

Overview

We adopt the supervised-learning approach to distinguish among EMR traces of different applications. The classifier is trained based on previously labeled traces, and then used to identify coming traces. We intercept slices of traces that match each application's operation time as the testing/training samples.

Feature Selection

Our classifier is trained on two feature types: the time-domain and the frequency-domain. For time-domain, the

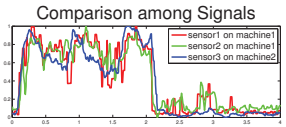

(a) App: Safari

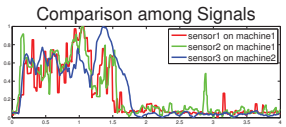

(b) App: Word

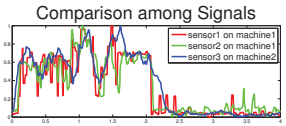

(c) App: Excel

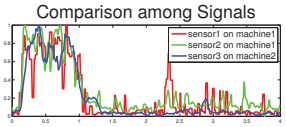

(d) App: iTunes

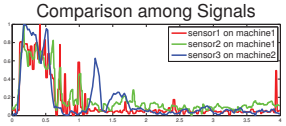

(e) App: VLC

Figure 3: Electromagnetic Traces Collected from Different Laptops using Different Magnetometers.

Learner ID	Features	Distance Metric
1	time-series	correlation coefficient
2	time-series	DTW cost
3	frequencies	correlation coefficient

Table 1: Learner Information.

EMR traces need to go through three preprocessing steps before training and classifying: normalization, cross-correlation and interpolation. For frequency-domain, we use FFT to extract their frequency distributions.

Training and Classification

We build three learners adopting the same nearest-centroid strategy but different classifying metrics, and use ensemble learning to generate the final label based on each learner's decision. First, we adopt the average DTW method [5] to generate a centroid for each class utilizing all the training samples that belong to that class. Then in each learner, we compare the coming testing sample with those centroids to find a centroid that is closest to that sample. The feature selection and distance metric of each learner is shown in Table 1. We use voting strategy to choose the label chosen by most of the learners. If three learners make contradictory decisions, the classifier will trust the most robust learner (learner 1).

Evaluation

We use three mobile phones with 3-axis magnetometer (Table 2) to collect the EMRs that belong to 10 applications on two Mac OS laptops. In general, the classification accuracy is up to 98% when the training/testing traces are collected from the same laptop using the same sensor, 98% when traces are collected from the same laptop using different sensors, and 61%when traces are collected from different laptops using different sensors, as shown in Figure 2.

Sensor ID	Model	Sampling Frequency
1	iphone SE	100Hz
2	Nexus 5	50Hz
3	iphone 5S	50Hz

Table 2: Sensor Information.

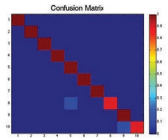

(a) Self Classification. Accuracy: 98%.

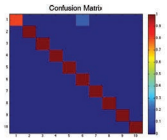

(b) Cross-sensor Classification. Accuracy: 98%.

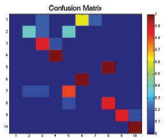

(c) Cross-machine Classification. Accuracy: 61%.

Figure 2: MagAttack Classification Performance.

Robust Features

We find that each application generates distinct EMR traces that are consistent through various sensor models or sampling rates. We conduct experiments in three cities: Hong Kong, Nanjing and Shenzhen respectively. We find the EMR traces to be independent of the earth's magnetic filed as well. Some of the traces are shown in Figure 3. The correlation matrix of these applications is shown in Figure 5, where the metric adopted is the correlation coefficient of time-domain features.

Impacts of Ensemble Learning

We train three learners and use voting strategy during classification so that the decision accuracy can be better than any single learner (Table 1). The performance comparison of each learner is shown in Figure 4. Note that in Figure 4 we collect EMR traces when other background processes are running so as to compare the robustness of each learner on noisy samples.

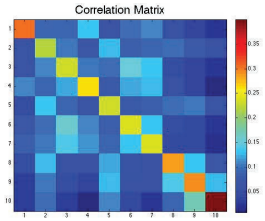

Figure 5: Correlation Matrix of Ten Apps: PowerPoint, Word, Excel, Chrome, Fire-Fox, Safari, Skype, iTunes, VLC, MPlayer.

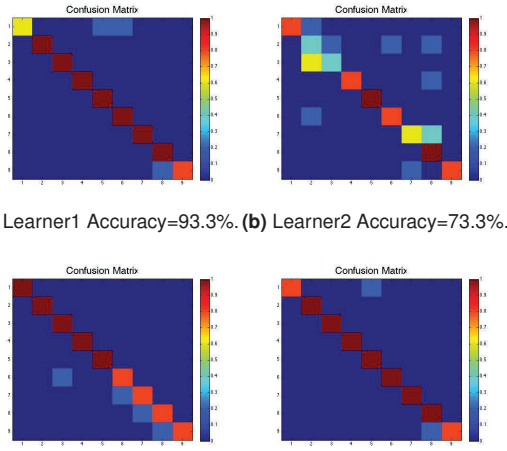

(a) Learner1 Accuracy=93.3%. **(b)** Learner2 Accuracy=73.3%.

(c) Learner3 Accuracy=91.1%. **(d)** Voting Accuracy=95.6%.

Figure 4: Comparison of Learner Performance.

Conclusion and Future Work

In this poster, we present MagAttack which uses a built-in magnetometer on a commodity mobile phone to sniff which App is running on a nearby computer by tracking changes in EMRs generated from the CPU of that computer. Our preliminary results show that MagAttack is independent of the attacker's mobile devices and the target's laptops, and can achieve a classification accuracy of up to 98%. Our future work will focus on decomposing EMR footprints while opening multiple applications at the same time as well as classifying webpages that the victim is browsing to explore more about the victim's habits and interests.

REFERENCES

1. Daniel Genkin, Itamar Pipman, and Eran Tromer. 2015. Get your hands off my laptop: Physical side-channel key-extraction attacks on PCs. *Journal of Cryptographic Engineering* 5, 2 (2015), 95–112.
2. Sidhant Gupta, Matthew S Reynolds, and Shwetak N Patel. 2010. ElectriSense: single-point sensing using EMI for electrical event detection and classification in the home. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*. ACM, 139–148.
3. Paul C Kocher. 1996. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *Advances in Cryptology-CRYPTO 1996*. Springer, 104–113.
4. Yan Michalevsky, Dan Boneh, and Gabi Nakibly. 2014. Gyrophone: Recognizing speech from gyroscope signals. In *23rd USENIX Security Symposium (USENIX Security 14)*. 1053–1067.
5. François Petitjean, Germain Forestier, Geoffrey I Webb, Ann E Nicholson, Yanping Chen, and Eamonn Keogh. 2014. Dynamic time warping averaging of time series allows faster and more accurate classification. In *Data Mining (ICDM), 2014 IEEE International Conference on*. IEEE, 470–479.
6. Cati Vaucelle, Hiroshi Ishii, and Joseph A Paradiso. 2009. Cost-effective wearable sensor to detect EMF. In *CHI'09 Extended Abstracts on Human Factors in Computing Systems*. ACM, 4309–4314.
7. Alenka Zajic and Milos Prvulovic. 2014. Experimental demonstration of electromagnetic information leakage from modern processor-memory systems. *Electromagnetic Compatibility, IEEE Transactions on* 56, 4 (2014), 885–893.