

Hard Drive Side-Channel Attacks using Smartphone Magnetic Field Sensors

Sebastian Biedermann¹, Stefan Katzenbeisser¹ and Jakub Szefer²

¹ Security Engineering Group, Technische Universität Darmstadt*
{biedermann, katzenbeisser}@seceng.informatik.tu-darmstadt.de

² Computer Architecture and Security Laboratory, Yale University
jakub.szefer@yale.edu

Abstract. In this paper we present a new class of side-channel attacks on computer hard drives. Hard drives contain one or more spinning disks made of a magnetic material. In addition, they contain different magnets which rapidly move the head to a target position on the disk to perform a write or a read. The magnetic fields from the disk's material and head are weak and well shielded. However, we show that the magnetic field due to the moving head can be picked up by sensors outside of the hard drive. With these measurements, we are able to deduce patterns about ongoing operations. For example, we can detect what type of the operating system is booting up or what application is being started. Most importantly, no special equipment is necessary. All attacks can be performed by using an unmodified smartphone placed in proximity of a hard drive.

1 Introduction

Hard drives are an integral part of almost any computer as they are the persistent storage medium for code and data. Operation of the hard drive is directly correlated with the type of workload being processed on the computer. The movement of the hard drive head is enabled by a magnetic field. This magnetic field can be picked up by sensors outside of the disk drive, and thus enables our new side-channel attacks. Whenever data is being written to, or read from, the head mechanism has to move, causing disturbance in the magnetic field which we can detect. Disk drives have previously been subject to research on covert timing channels, e.g. [3]. However, these attacks required direct access to the target computer. The magnetic side-channel, on the other hand, is a non-invasive attack which can be carried out without physical contact. This attack falls under the broad class of electromagnetic (EM) attacks. EM attacks have been carried out on various computer components, such as CMOS chips [5], or smart cards [1]. Disk drives, however, have not been subject to such attacks.

* This work was supported by CASED (www.cased.de), EC-SPRIDE (www.ecspride.de) and the the German Federal Ministry of Education and Research (BMBF) under grant no. 01C12S01V.

Our research was motivated by two trends in industry. First, there has been a great proliferation of smartphones with various sensors. The key sensor which we focus on is the digital compass, also called the magnetometer. Today’s phones have sensors which can pick up magnetic fields having a strength of μT (micro Tesla, where Tesla is the unit of magnetic field strength or magnetic flux density). Second, the industry is constantly working on optimizing the size of the servers and computer equipment. This has led to smaller and thinner computers. The reduced size means there is less shielding and shorter distance from the disk drive to the outside of the computer chassis. This creates a situation where magnetic field fluctuations can be more easily detected using sensors such as those available on today’s phones. Attacks are performed by using an application running on an unmodified smartphone. Our side-channel attacks allow us to:

- detect the operating system that is used;
- distinguish between known applications being started;
- distinguish Virtual Machine activity on a server;
- match ongoing network traffic to a server; and
- detect file caching based on disk activity.

2 Magnetic Field and Hard Drives

In this section we provide a short summary on magnetic fields and magnetic disk drives. More information is available in a variety of books, such as [7].

2.1 Magnetic Fields

Tesla (T) is the unit of magnetic field strength, often denoted as B . It can also be expressed in units of gauss (G), where 1 gauss equals 100 μT . As a point of reference, Earth’s magnetic field ranges between 50 to 75 μT . The magnetic field strength in close proximity of a hard drive can vary due to the disk operation and the magnets inside the disk drive. In particular, the movement of mechanisms inside the hard drive when accessing data causes the magnetic field strength to fluctuate by about 3 μT when sensed next to hard drive ($d \approx 0$ cm). The field strength B decreases as a square of the distance d with $B \propto \frac{1}{d^2}$ from a point source. Given the magnetic field strength B_1 at distance d_1 , the strength B_2 at distance d_2 can be inferred from $\frac{B_1}{B_2} \propto \frac{d_2^2}{d_1^2}$. Background magnetic field fluctuation (i.e. noise) always is about 0.1 μT in magnitude in our experiments. Thus, realistic maximum measurement distances d_i are constrained by the fact that the field changes need to be above the background noise, thus $B_i > 0.1 \mu\text{T}$.

2.2 Hard Drive Magnetic Fields

There are three sources of magnetic fields in a standard hard drive: a) magnetic disk platters, b) the disk drive head, and c) the mechanisms for moving the disk drive head. Both a) and b) are too weak to be detected outside a disk drive

chassis. We show, however, that detecting the strength of c) is feasible. In a hard drive, there are two magnets, one above and one below the head movement assembly. The head movement assembly includes a coil of wire. When current is passed through the wire, a magnetic field is generated, which causes the head assembly to displace. Depending on the direction of the current, the generated field interacts with the fixed magnets and causes the head movement assembly to move left or right from its rest position. At rest the head assembly is either all the way to the left or all the way to the right, depending on the particular disk in use. The strength of the field determines how far the head is displaced.

2.3 Magnetic Field Sensors on Mobile Devices

On a modern smartphone the magnetic field sensor measures the strength of the magnetic field along three axes. On Android-based devices, applications may access the `TYPE_MAGNETIC_FIELD` sensor to get the field readings. The sensor outputs data on the strength of magnetic field in units of μT along three axes. We experimentally verify that when placing the phone next to a drive, the z axis measurements give the least noisy readings. If the phone is placed on top of the drive, the x or y axis measurements provide best results assuming minimal curvature with the z axis perpendicular to both. Note, however, that the earth's magnetic field is parallel to most surfaces (x or y axis when the phone is laying flat on some surface) which can cause high noise along these axis. Thus, ideally the axis perpendicular to the earth's surface should be used for the measurements, which need not always be the z -axis.

2.4 Magnetic Field Measurements

In our experiments we use unmodified Samsung Galaxy S4 Mini and Samsung Galaxy S2 smartphones. We access their magnetometer using a custom application, however one which requires no special permissions (unlike applications that may require permissions to use camera or other sensors). Before any experiment is performed, the application measures the background magnetic field in order to calibrate subsequent measurements. This is done by taking 100 measurements along the x , y and z axis and computing the mean of all measurements. This way the average field strength $(\bar{x}, \bar{y}, \bar{z})$ of the background noise can be established. Once the background magnetic field strength is measured, the application records measurements M_i along the x , y and z axis which correspond to the difference between the newly measured value and the average background noise, $M_i = (\text{timestamp}_i, x_i - \bar{x}, y_i - \bar{y}, z_i - \bar{z})$.

3 Magnetic Side-Channel Attacks

This section presents a variety of magnetic side-channel attacks which we have explored. We focus on two targets. First, we launch attacks against a laptop in an office-type environment. Second, we launch attacks against a server co-located in a server rack with other running servers.

3.1 Attack Composition

The activity of the hard drive can be matched to changes in the magnetic field strength, since the head is moved each time files are written or read on different locations. Additionally, the magnitude of the magnetic field strength can be matched to different locations on a hard drive on which the head currently operates. These two different factors create characteristic fingerprints over time that can be matched to ongoing operations of a hard drive. The smartphone should be located closely, at best in a distance of 3 – 4 cm to the target hard drive. In order to detect certain operations on an arbitrary target hard drive, we analyze correlations between measurements taken in an enrollment phase and attack measurements. During enrollment, we record several measurements of the same hard drive operation and compute the average of all measurements, which yields to a characteristic enrollment vector for one specific activity. Before computing the average, we synchronize all individual measurements. We do this by shifting the measured vectors by ± 0.5 seconds; the shift which produces the best correlation with the previous measurements is taken. In the attack phase, we compute the correlation between a new measurement and all stored enrollment vectors. Again, in order to synchronize the new measurement with the enrollment vectors, we shift it by ± 0.5 seconds. Finally, we classify the measurement to the enrollment class which achieves the highest correlation.

3.2 Example Attacks against Laptop

The target laptop is an Acer Aspire 5733z with a Toshiba 320 GB disk drive with 5400 rpm. The attack setup is shown in Figure 1. The smartphone is placed in front of the laptop, near where the disk drive is mounted.

As one possible scenario, we envision a malware-infested smartphone. The malware, which the user downloaded when searching for a “digital compass” application, is running in the background on the phone and collecting magnetic




Fig. 1. Laptop attack setup, shown with the Samsung Galaxy S4 Mini.

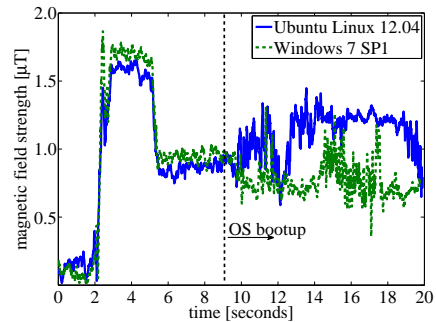


Fig. 2. Field strength while booting Linux and Win7 (average of 20 runs).

sensor readings once it is triggered. The unsuspecting user may place the phone in close proximity to the laptop when working on the laptop in the office or in a cafe. In this position, the malware can launch a number of attacks, such as OS or application startup detection.

OS Boot-Up Detection: First, we use the measurements to investigate which operating system (OS) is booted on the laptop. Figure 2 shows recorded measurements during the boot-up of Ubuntu Linux 12.04 (64 Bit) and Windows 7 SP1 (64 Bit) on the same laptop, each taken for 20 seconds right after turning on the notebook. The first nine seconds are very similar in both curves, since first the BIOS is loaded and the hard drive is initialized. After nine seconds, the OS starts booting and differences in the course of the magnetic field can clearly be seen. These characteristic deviations are based on the different underlying file system as well as on different processes which are starting during boot-up. During enrollment we record vectors with 1000 measurements (ten seconds) starting nine seconds after the first change in the magnetic field could be detected. This way, we skip the BIOS and the firmware and only use the characteristic measurements of the operating system’s boot-up. The enrollment vectors are averages over ten independent trials. Table 1 shows the average Pearson correlation between ten measurement vectors taken from different boot-ups of different OSs and enrollment vectors for Ubuntu Linux, Windows 7 and Windows Vista. In each case, the strongest correlation occurs between the measurement vector and the enrollment vector of the same OS. The correlation to other enrollment vectors of other OSs is small in comparison. We can use the correlation between attack measurements and enrollment measurements in order to decide which OS was booting. In this experiment, we use five measurements of each OS boot-up procedure for enrollment and ten new measurements which we classify to one of the three OS classes depending on their highest correlation. The results of this simple classification approach lead to an average error rate of 3% across 20 enrollments and 20 attack measurements in each OS setup with which an attack vector was not correctly classified.

Application Start-Up Detection: In the next attack, we use the changes in the strength of the magnetic field to detect the start of some known applications. We use the same setup as before. The laptop runs Windows 7 SP1 (64 Bit) and has three different browsers, namely Microsoft Internet Explorer (ver. 8), Google Chrome (ver. 34) and Mozilla Firefox (ver. 28). Our attack goal is

measurement	enrollment vectors		
	Linux	Win7	WinVista
Linux	0.40 ± 0.22	0.09 ± 0.08	0.12 ± 0.10
Win7	0.08 ± 0.04	0.29 ± 0.10	0.14 ± 0.06
WinVista	0.04 ± 0.13	0.06 ± 0.12	0.37 ± 0.22

Table 1. Average correlation between new attack measurements recorded during booting different operating systems and the enrollment vectors.

	enrollment vectors		
measurement	IE	Chrome	Firefox
IE	0.34 \pm 0.11	0.23 \pm 0.10	0.23 \pm 0.08
Chrome	0.29 \pm 0.14	0.47 \pm 0.11	0.07 \pm 0.19
Firefox	0.22 \pm 0.13	0.13 \pm 0.15	0.49 \pm 0.16

Table 2. Average correlation between new attack measurements recorded during start of different browsers and the corresponding enrollment vectors.

to determine which browser was started. Table 2 shows the average Pearson correlation between measurements and enrollment vectors recorded during the start of IE, Chrome and Firefox. In each case, the strongest correlation occurs between the measurements and the enrollment vectors of the same browser.

Based on the highest correlation, we can classify new samples to a browser class with an average error rate of 23%. The results indicate that the start-up of different applications can be distinguished with a good accuracy.

3.3 Example Attacks against Servers

In this section we present attacks against a server co-located in a server rack with other servers. The target server is a Dell R210 with a 2TB (7.2K RPM SATA) disk drive. The smartphone is placed on top of the server, near where the disk drive is mounted. In order to successfully mount the attack and to be able to place the smartphone, at least one rack above the target server needs to be empty. This also results in a distance of at least $1\frac{3}{4}$ inches (1 Rack Unit) between our target server and the next server. The distance greatly decreases the influence of other operating hard drives in the same server rack. As one possible attack scenario, we envision an unscrupulous data center employee who wants to gain information about processes running on the servers.

VM Activity Detection: In this example, an attacker tries to gather insights about virtual machines (VMs). Our server runs the Xen hypervisor 4.2 and different VMs, each having a storage of 128 GB. We measure the magnetic field during writing a file in different VMs vm01, vm02 and vm03. Figure 3 shows the average magnetic field measurements. The average strength can be used to estimate which VM is currently operating on the server’s hard drive. Given an attack measurement, we again compute the average magnetic field strength and perform a classification based on the smallest difference to an enrollment measurement. Results showed that the activity of the server’s hard drive can be assigned to vm01 or vm03 with an average error rate of 35%. The differences between the magnetic field strengths of the VMs are small, but detectable.

Host Server Detection: As another example, the attacker may know that a website is hosted on a server in a data center and she wants to find out exactly which server actually hosts the site. To reach this goal, the attacker triggers downloads from the website with the smartphone while measuring the magnetic field radiated by servers in the rack in multiple trials. This way, the attacker can determine if a server created a magnetic field corresponding to a read operation

of a data block having the appropriate size that matches the download. In order to test the effectiveness of our attack, we use ten enrollment measurements recorded during downloading a 32 MB file from our Web server and computed and average enrollment vector. We found that the attack and enrollment measurements correlate with 0.33 ± 0.11 if the file is downloaded from the server on which the smartphone is located, while the attack measurements correlate only with 0.12 ± 0.07 if the file is not downloaded from this server. Classifying 20 new attack measurements based on the highest correlation to the enrollment vector while only during 10 of the measurements the file is downloaded lead to an average error rate of 15%. The correlation can be used to reveal the server that hosts the website which provides the download.

File Caching Detection: If a server caches files in memory, then host server detection attack can be prevented. However, detection of the caching behavior is also an interesting attack objective. Frequently accessed files can be distinguished from infrequently accessed ones – thus leaking information whether a file has recently been in use. Figure 4 shows disk activity (darker red regions show a larger change in the magnetic field over time) when several files of 16 MB are accessed, in this case downloaded from the Web server. Two of the files (namely 5 and 11) have been accessed before the test and became cached. To test the attack, we perform 10 enrollment measurements while downloading files and 10 enrollment measurements without downloading. Subsequently, we perform 20 attack measurements while downloading 20 different files, 10 of them were already downloaded before and are cached, 10 of them were never downloaded. We classify the attack measurements to one of the two classes “is read from hard drive” or “is not read from hard drive”, the latter means the file was cached in memory. We do this by choosing the class that yields to the highest correlation. According to our results, an attacker can distinguish if a file has already been accessed or not on a Web server with an average error rate of 5%.




Fig. 3. Strength during writing data in vm01, vm02 and vm03 (50 runs).




Fig. 4. Field while downloading files; 05 and 11 have already been accessed.

4 Related Work

There are a lot of other side-channel attacks, but usually they require access to the victim system or physical connection to measure the power usage, while our attack requires no physical connection, but only physical proximity. In 1984, Gold et al. [3] presented analysis of covert channels due to placement of the disk arm. While by 1991, Karger et al. [4] presented research on storage channels due to hard disk drive head movement. However, the majority of the past work on electromagnetic (EM) side-channels has focused on processors. Researchers used specialized magnetic sensors to sense emanations and recover a secret key [5]. Other work has shown that the electromagnetic attack on processors can obtain at least as much information as power consumption based side-channels [6]. Given the need for proximity when working with electromagnetic emanations, research has focused on smart cards where physical access and proximity are easy. Researchers have shown electromagnetic side-channel attacks on various smart cards with different hardware protections, and still were able to recover the encryption keys [2]. Others were even able to propose a model that completely and quantitatively expresses the information leaked from electromagnetic side-channel in CMOS devices, such as smart cards [1]. Purely using magnetic field measurements against hard drives has not been explored yet.

5 Conclusion

In this paper we presented a new class of side-channel attacks on computer hard drives. From measurements of the magnetic field, which carries information about the movement of the hard drive mechanisms, we are able to deduce patterns about ongoing operations. All experiments were performed using a modern, unmodified smartphone which was placed in proximity to a hard drive, even outside a laptop or a server chassis.

References

1. D. Agrawal, B. Archambeault, J. Rao, and P. Rohatgi. The em side channel(s). In *Cryptographic Hardware and Embedded Systems*, volume 2523, pages 29–45. 2003.
2. K. Gandolfi, C. Mourtel, and F. Olivier. Electromagnetic analysis: Concrete results. In *Cryptographic Hardware and Embedded Systems*, pages 251–261. 2001.
3. B. D. Gold, R. Linde, and P. F. Cudney. Kvm/370 in retrospect. In *IEEE Symposium on Security and Privacy*, pages 13–23, 1984.
4. P. A. Karger and J. C. Wray. Storage channels in disk arm optimization. In *IEEE Symposium on Security and Privacy*, pages 52–63, 1991.
5. E. Mateos and C. Gebotys. Side channel analysis using giant magneto-resistive (gmr) sensors. In *COSADE workshop proceedings published by CASED*, 2011.
6. J.-J. Quisquater and D. Samyde. Electromagnetic analysis (ema): Measures and counter-measures for smart cards. In *Proceedings of the Int. Conference on Research in Smart Cards: Smart Card Programming and Security*, pages 200–210, 2001.
7. N. Rao. *Fundamentals of Electromagnetics for Electrical and Computer Engineering*. Pearson Education, 2011.