EMLogger: Inferring Computer Activities via EM Side-channel of Disks

Wenfan Song[1], Jianwei Liu[1, 2], Jinsong Han ✉[1]

[1]Zhejiang University, Hangzhou, China.
[2]Hangzhou City University, Hangzhou, China
{wenfansong, jianweiliu, hanjinsong}@zju.edu.cn.

Abstract

In computer systems, the built-in disk is essential for data storage as computers typically need to read and write data from the disk during operations. Some existing works have utilized the read/write characteristics of disks to infer sensitive information, such as browsing websites. However, previous research primarily focuses on hard disk drives (HDDs). With solid state drives (SSDs) gradually becoming the mainstream option for computer disks, these approaches have shown limitations. In this paper, we reveal a novel side-channel vulnerability targeting both HDDs and SSDs, named EMLogger. Specifically, we find that attackers can exploit the electromagnetic (EM) radiation leaked by the disk to detect ongoing activities on the computer. To enhance the strength of EM signals, we propose a sub-signal fusion method. Besides, we employ machine learning techniques for feature extraction and activity classification from the enhanced EM signals. Finally, we conduct real-world experiments on computers equipped with HDDs or SSDs. Our experimental results demonstrate that EMLogger achieves an accuracy of over 98% in inferring computer activities. Furthermore, the experiments validate the robustness of EMLogger at varying attack distances.

Keywords: Electromagnetic Emanation, Side-channel Attack, Disk.

1 Introduction

The built-in disks, including hard disk drives (HDDs) and solid state drives (SSDs), are essential in computer systems for storing operating systems, applications, and user data. However, recent studies [1, 2] have demonstrated that disks can be exploited as an attack vector to steal sensitive information from the host, such as the websites being accessed. This is due to the dependency of computer activities, like browsing

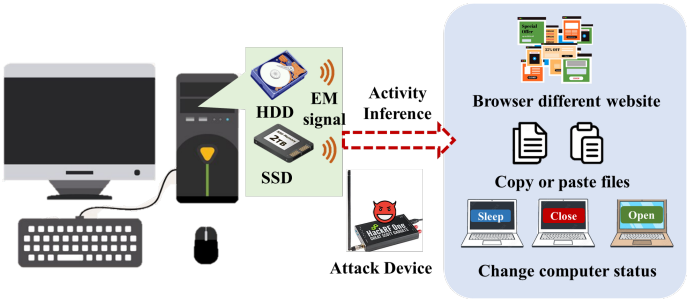

Fig. 1 EMLogger can capture EM signals leaked from the computer's disk through an attack device to infer computer activities, thereby endangering the user privacy.

websites and watching videos, on the data read and write operations performed by the disk. Specifically, these studies have focused on leveraging the physical movements of the read/write head in HDDs to infer ongoing activities on the host [1, 2]. The read-/write head's position and movements over the disk's platters can reveal the computer activities. However, these approaches [1, 2] have a limitation: they are only applicable to HDDs, whereas a majority of computers now use SSDs [3]. This is because SSDs are based on flash storage technology, and do not rely on the mechanical movement of magnetic heads and rotating platters like HDDs.

In this paper, we introduce a novel physical side-channel vulnerability, namely EMLogger, targeting both HDDs and SSDs to infer computer activities. Specifically, we find that the EM signals leaked by the common components in both HDDs and SSDs, i.e., the clock modules and DRAM modules, can reveal fine-grained information about the computer's ongoing activities. As shown in Fig 1, an attacker can hide the attack device under a table, When the victim operates the computer on this table, the attack device can capture the EM signals emitted by the disk inside the computer. The attacker can then infer the host activities based on these EM signals. This poses a serious threat to the victim's privacy. For instance, an attacker could determine which websites the victim is browsing. The leakage of website information provides attackers with at least two key insights. First, attackers can target precise advertisements based on the victim's preferences. Second, attackers can infer sensitive information about the victim, such as their interests and health status.

To implement EMlogger, we address two challenges. **(1) How to capture high-quality EM signals in a noisy environment?** On one hand, electromagnetic compatibility (EMC) standards impose strict limits on the intensity of the EM emanation emitted by the disk [4]. On the other hand, other electronic devices in the ambient environment also emit EM signals, which may interfere with the EM signal of the disk. These factors result in a low signal-to-noise ratio (SNR) for the captured EM signals from the disk. To tackle this challenge, we adopt a sub-signal fusion technique. Particularly, the EM radiation from the disk consists of a series of sub-clock signals. The amplitude of these sub-clocks varies consistently over time, and the frequency intervals between these sub-clocks are uniform. Based on these consistencies, we aggregate the sub-clock signals in the frequency domain to enhance the leaked EM signals from the disk. **(2) How to extract representative features from the EM signals emitted by the disk to infer computer activities?** We first theoretically

analyze and model the EM signals leaked by the disk to explore characteristics related to computer activities. Then, we use machine learning to further learn and classify these features to accurately infer computer activities. To validate the threat posed by EMLogger, we also conduct real-world attack experiments on computers equipped with HDDs and SSDs. The experimental results indicate that the accuracy of inferring computer activities can reach over 98%. Additionally, we confirm the robustness of EMLogger at different attack distances: within a range of 15 cm, the accuracy remains above 83%. In summary, our contributions are as follows:

- We explore a new electromagnetic side-channel attack targeting both HDDs and SSDs, namely EMLogger. Our theoretical model demonstrates that activities on a computer can be covertly monitored through the EM side channel.
- We propose a sub-signal fusion method to enhance the SNR of the EM signals leaked by the disk. Besides, we also adopt machine learning techniques to extract activity-related features.
- We conduct real-world experiments on computers equipped with HDDs and SSDs. The experimental results show that EMLogger can achieve an accuracy of over 98% in inferring computer activities, which demonstrates the potential privacy threats posed by EMLogger.

2 Related Work

EM emanations leaked by computing devices have been widely leveraged for various side-channel attacks, such as theft of encryption keys [5, 6] and interference with private information [1, 7–9]. These emanations are primarily generated by the high-speed switching circuits [9] and clock modules [5–7] within the devices, which makes them valuable for constructing covert channels for data leakage. The essence of using EM signals for side-channel attacks is that different activities inside the device emit EM signals with distinct characteristics. For example, Alam *et al.* [5] and Genkin *et al.* [6] successfully recover the full RSA key and ECDSA secret signing keys via EM emanations, respectively. The work in [7] infers website fingerprinting and keystroke timing vulnerabilities in GPUs using EM side-channel analysis. Moreover, Enev *et al.* [9] use in-display fingerprint sensors to extract fingerprint information for deceiving and unlocking smartphones. The work in [8] utilizes EM emanations from in-display fingerprint sensors to extract fingerprint information for deceiving and unlocking smartphones. Biedermann*et al.* [1] exploit the EM emanations from HDD heads for attacks. Similar to this work, we use EM signals leaked from the disk to infer the computer activities. Despite the similarities, our approach diverges fundamentally and demonstrates superior efficacy. Firstly, our attack has a broader scope, encompassing both HDDs and SSDs, whereas work [1] is limited to HDDs. This difference arises from our utilization of EM signals emitted by common components of HDDs and SSDs (i.e., clock module and DRAM module) to achieve attacks, while the work [1] relies on the unique EM signals emitted by HDD magnetic heads. Secondly, the attack distance of our method is longer (15 cm) compared to that of [1] (3 cm). This is because we have investigated techniques to amplify the EM signals of interest, enabling effective EM signals capture at larger distances. Finally, the accuracy of our approach far exceeds

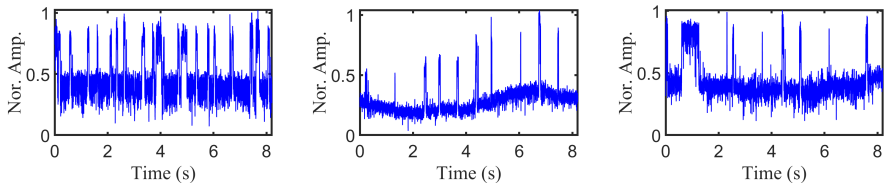

(a) eBay. (b) Tiktok. (c) Google.

Fig. 2 EM signals leaked from the disk when browsing different websites. Nor. Amp means normalized amplitude.

that of [1] (i.e., 98% » 75%), attributed to our use of deep neural networks (DNN) to effectively extract fine-grained features from EM signals, which leads to efficient classification.

3 Threat Model

The goal of EMLogger is to eavesdrop on the EM signals emitted by the disk inside the victim computer and then extract sensitive information, such as the websites the victim is browsing. In real-world scenarios, we assume the victim places his/her computer (e.g., a laptop) on a table in a public place (e.g., a cafe or meeting room). The attacker pre-hides the attack device (i.e., a software-defined radio (SDR) like a HackRF) near the victim computer (such as under the table). In this case, the attacker can remotely monitor the EM leakage from the disk inside the victim computer. Since the activities of the computer rely on the disk, the EM emanations from the disk can be utilized to infer the computer's activities. Subsequently, the attacker can obtain the user's sensitive information, such as hobbies, and behavioral habits. Furthermore, we assume that attackers can determine the computer model by observing its appearance, and then further search online for the model of the disk inside the computer. Subsequently, the attacker can acquire an identical model of the computer from the market for further study. It is worth noting that the victim computer is not assumed to have any malicious software or physical tampering.

4 Privacy Leakage Exploration

To investigate the feasibility of EMLogger attack, we first develop a theoretical model for the EM emanations of hard drives and the associated privacy leakage process. Subsequently, we conduct a series of experiments to demonstrate the viability of such EM side-channel attacks.

4.1 Side-channel Modeling

Different activities performed on a computer exert various loads on its inside disk, which depends on the frequency and volume of disk reads and writes. The fluctuation in disk load induces changes in the EM signals emitted by the disk. Consequently, by eavesdropping on these EM signals, we can infer the computer's activities.

The EM emanations of the disk mainly originate from two modules: the clock module and the DRAM module. The clock module is responsible for coordinating the

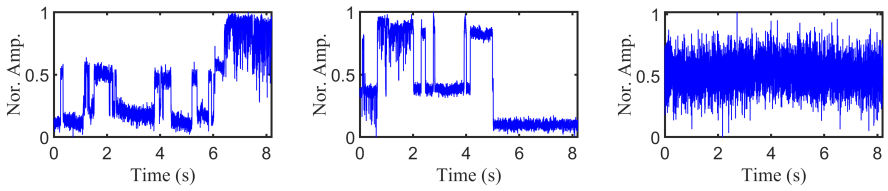

(a) Power-on state. (b) Power-off state. (c) Sleep state.

Fig. 3 EM signals leaked from the disk when the computer is in different power states. Nor. Amp means normalized amplitude.

timing sequence of read and write operations and controlling the data transfer rate. Meanwhile, the DRAM module is used for temporary storage and quick access to data and instructions being processed. For the clock module, we can observe that its EM signals $m_{clk}(t)$ consist of a series of sub-clocks, usually distributed on the side-band below the clock frequency f_0, which can be expressed as:

$$m_{clk}(t) = \{m_{clk}[1,t], m_{clk}[2,t], ..., m_{clk}[N,t]\}. \tag{1}$$

Specifically, the n-th sub-clock can be expressed as:

$$m_{clk}[n,t] = A_{clk}(n)\sin(2\pi f_{clk,n}t), n \in [1,N] \tag{2}$$

where $A_{clk}(n)$ and $f_{clk,n} = f_0 - nf_m, 1 \leq n \leq N$ are the amplitude and frequency of the n-th sub-clock, respectively. f_m represents the frequency interval between adjacent sub-clocks, and N represents the total number of sub-clocks. When the computer conducts activities (e.g., browsing websites), it performs a series of reading and writing operations on the DRAM module of the disk. This leads to the emission of EM signals by the disk (denoted as $m_{dram}(t)$), which are then amplitude-modulated onto the clock's electromagnetic signal ($m_{clk}(t)$). It is worth noting that although $m_{dram}(t)$ alters the amplitude of $m_{clk}(t)$, it does not affect the spectral pattern of the clock. Considering the effect of host activity on the load of the DRAM (denoted as $\alpha(t)$), the EM signal of the DRAM module can be further expressed as $m_{dram}\{\alpha(t)\}$. In this case, the EM emanation leaked by the disk (denoted as $m_{disk}(t)$) can be expressed as:

$$m_{disk}(t) = m_{dram}\{\alpha(t)\} * \{m_{clk}[1,t], m_{clk}[2,t], ..., m_{clk}[N,t]\}. \tag{3}$$

In this formula, we can observe that the intensity of the sub-clock's electromagnetic radiation varies in proportion to the load of the (i.e., $\alpha(t)$). Due to the strong correlation between $\alpha(t)$ and host activities, different activities introduce different DRAM loads. Therefore, by analyzing the EM signals leaked by the disk, an attacker can infer the host activities, such as browsing different websites.

4.2 Preliminary Experiment Validation

In this work, we focus on three kinds of activities that could leak user privacy, namely (1) browsing different websites, (2) the power state of the computer, and (3) writing to/reading from files to the disk. To explore the potential for identifying these activities

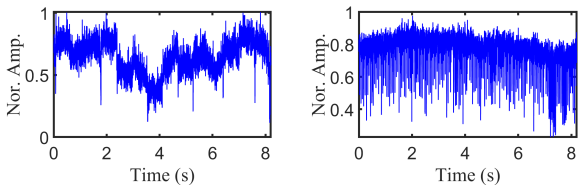

(a) Writing files. (b) Reading files.

Fig. 4 EM signals leaked from the disk when writing/reading files to/from the disk.

via EM signals, we conduct a series of preliminary experiments. Specifically, we use a software-defined radio (HackRF SDR), a Foresight low-noise amplifier, and a 3dBi antenna to collect the EM signals emitted by the internal disk of a Lenovo Xiaoxin Pro 14 laptop while performing different activities. Since the amplitude variation of sub-clocks over time is similar (determined by the DRAM load $\alpha(t)$), we only capture one sub-clock signal to explore inferring computer activities using EM signals.

■ **Websites.** In this experiment, we browse eBay, Ticktock, and Google respectively, and record the corresponding EM emanations. Figure 2 shows the normalized EM signals while the computer performs various activities. We can observe that there are noticeable differences in the EM signals corresponding to these three websites. This is because each website caches different content to the disk, leading to variations in the disk load. Therefore, the EM signals corresponding to different websites exhibit distinct characteristics. Consequently, attackers can exploit these EM signals to infer the websites that users are browsing.

■ **Power state.** Similar to the website inference experiment, we set the computer to power-on, power-off, and sleep states separately and record the corresponding EM signals leaked from the disk. Figure 3 shows significant differences in the EM signals corresponding to these three computer power states. This is because the load of the disk varies depending on the power state. In sleep mode, the disk is in a low-power state. When powered on, the disk transitions from a powered-off state to an active state, and vice versa when powered off. Therefore, attackers can utilize EM signals to detect the computer's power state.

■ **Write/read files to/from the computer.** In this experiment, we write/read files to/from the computer and record the corresponding EM signals. As shown in Fig 4, we can observe significant differences in the EM signals corresponding to these two types of computer activities. The reason is that these two types of activities involve various hardware operations and data transfer processes, resulting in differences in the load of the disk. Therefore, attackers can exploit EM signals to detect the read/write activities of the computer.

5 Methodology

With the captured EM signals, an attacker can infer the activities being performed by a computer. To achieve this goal, we first employ sub-clock fusion to enhance the signal-to-noise ratio (SNR) of the EM signals. After that, we induce a DNN-based approach to accurately infer computer activities.

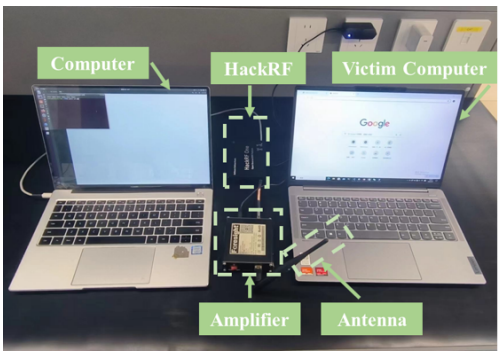

Fig. 5 EMLogger setup.

5.1 Sub-clock Fusion

As mentioned in Sec.4.1, the amplitudes of the sub-clock signals can reflect the variation in disk load (i.e., $\alpha(t)$), which depends on the activities being executed by the computer. Nevertheless, unlike the clock signals which are distributed at a single frequency, the EM emanations from the disk clock module are spread across a range of frequencies. This dispersion significantly reduces the signal-to-noise ratio (SNR) of the clock EM signal [10]. Additionally, electromagnetic noise from surrounding devices further decreases the SNR of the disk EM signals. To address the challenge of low SNR, we adopt the sub-clock fusion technique [7] to amplify the EM signal of the disk.

The core idea of sub-clock fusion is that (1) the amplitude variations of different sub-clocks are consistent over time, and (2) the frequency intervals between adjacent sub-clocks are uniform. This consistency allows us to effectively enhance the SNR of the clock module's EM signals by summing these sub-clock signals. Specifically, assume that the EM signal of the clock module (i.e., m_{clk}) comprises N sub-clocks, where $m_{clk}[i,t](i \in [1,N])$ represents the i-th sub-clock signal. We can obtain the spectrum of these N sub-clock signals through the short-time Fourier transform. Subsequently, we sum these N sub-clocks and obtain the fused clock signal (denoted as $m_{fused}(t)$) ,which can be expressed as:

$$m_{fused}(t) = \sum_{n=1}^{N} A_{clk}(n)\sin(2\pi f_{clk,n}t). \tag{4}$$

We can observe that the amplitude of the fused clock signal (i.e., $|m_{fused}(t)|$) is significantly enhanced compared to that of the i-th sub-clock ($|A_{clk}(i)\sin(2\pi f_{clk,i}t)|$). In this case, the amplified EM signal of the disk can be expressed as $m'_{disk} = m_{fused}(t) * m_{dram}\{\alpha(t)\}$. In the following, we will utilize a DNN-based approach to extract features of the disk load ($i.e., \alpha(t)$) from the amplitude trace of m'_{disk}, thereby accurately inferring the activities of the computer.

5.2 Activity Inference

Given the excellent feature extraction capabilities of DNN [11], we attempt to design a DNN-based learning model to mine the 'deep-hidden' features of the disk load from the amplitude trace of m'_{disk}.

As the amplitude traces are time-series data and one-dimensional convolution kernels are efficient in feature extraction, we opt for one-dimensional convolutional layers as the primary components of the DNN. To extract deep-level disk load features for interring computer activities, we employ three convolutional layers. Each convolutional kernel involves tuning three parameters: channel number, kernel size, and stride. The channel number critically influences the model's representational capacity: too few channels risk information loss, whereas an excessive number may escalate computational demands. Similarly, the kernel size dictates the granularity in feature extraction; too diminutive a kernel might overlook crucial features, while a large one could yield overly generalized features. The stride parameter governs the size of output feature maps, with an excessively small stride risking information loss and a too-large one imposing higher computational costs and memory usage. Based on our experiments, we set the kernel size to $16*1$ and the stride to $8*1$. The channel numbers of the three convolutional layers are set to 16, 32, and 64, respectively. Additionally, each convolutional layer is followed by a batch normalization (BN) function [12] and a rectified linear unit (ReLU) [13], which help enhance the network's generalization ability and improve its expressive power. Subsequently, we add two fully connected layers after the third convolutional layer to map the features into a probability space. To enhance the DNN's nonlinearity, we also add a sigmoid function after the fully connected layers. Then, we employ the cross-entropy loss function to measure the difference between the probability distribution output by the model and the actual labels, which continuously optimizes the trained DNN model. The cross-entropy loss function can be expressed as:

$$Loss = \frac{1}{N}\sum_{i=1}^{N} q_i log(p_i), \tag{5}$$

where N represents the number of the activity categories. q_i represents the probability distribution of the true activity category labels. p_i represents the probability distribution of the activity categories predicted by the model. Based on computational loss, we utilize the Adam optimizer to iteratively update the parameters in the DNN until the loss converges or reaches a predefined stopping condition.

6 Evaluation

In this section, we first introduce the real-world experiment setup and then detail the performance of EMLogger attacks.

■ **Experiment Setup.** We build the prototype of EMLogger based on a software-defined radio (HackRF SDR), a Foresight low-noise amplifier, and a 3dBi omnidirectional antenna. As shown in Fig 5, we place the antenna above the internal disk of the compromised disk to capture its leaked EM signals. These received signals are then amplified by the amplifier before being transmitted to the HackRF. Following this,

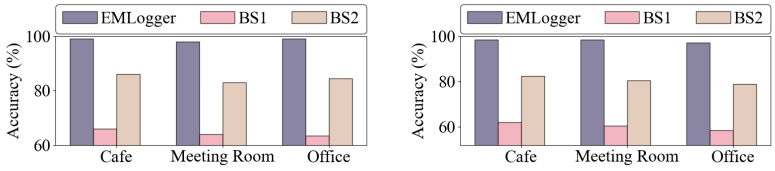

(a) Accuracy of the HDD. (b) Accuracy of the SSD.

Fig. 6 Activity classification accuracy of HDD and SSD under different environments.

we employ a Huawei KLV-W19L PC equipped with an Intel(R) Core(TM) i5-8265U CPU to further process the EM signals received from the HackRF.

■ **Data Collection.** To evaluate the robustness of EMLogger across various environments,we conduct experiments in three environments: a cafe, an office, and a meeting room. In each environment, we test the performance of EMLogger using two types of internal disks within the host: the HDD (model WD10EZEX) within the computer (model Dell OptiPlex 7060) and the SSD (model UMIS AM6A1) within the computer (model Lenovo Xiaoxin 14pro). For each disk, we capture its EM signals when the computer is executing 10 common activities. These ten activities are divided into three major categories: (1) web browsing, including eBay, Ticktock, BiliBili, Spotify, and Google; (2) computer power status: including power-on, power-off, and sleep states; (3) file transfer: copying or moving files to the disk. EM signals are sampled at a rate of 20 MHz with a bandwidth of 10 MHz. We collect over 4000 EM signals, with each EM signal corresponding to one activity the computer performed. Each EM signal comprises 5000 samples within a 4.1-second duration.

■ **Metric.** We define accuracy for computer activity class determination. The accuracy is the probability that an activity class of an EM signal sample is correctly identified.

6.1 Overall Effectiveness

We first evaluate the overall performance of EMLogger in three different environments. Particularly, we use 50% signal samples for DNN model training and the rest 50% for testing. To show the superiority of our sub-clock fusion method and the DNN method, we compare EMLogger with two baselines. Baseline (BS1): we directly use DNN to classify the raw EM signals without signal enhancement. Baseline (BS2): we use the support vector classifier (SVM) to classify the enhanced EM signal. Figure 6 shows the results of accuracy of inferring activities from EM signals leaked by the HDD and SSD under three different environments. We can observe that the average accuracy of EMLogger exceeds 98%, while the averages of these two baselines are 62.4% and 82.6%, respectively. The high accuracy of EMLogger demonstrates the outstanding classification performance in different environments. These comparisons with two baselines indicate that sub-clock fusion and DNN classification can effectively enhance the EM signals and extract features.

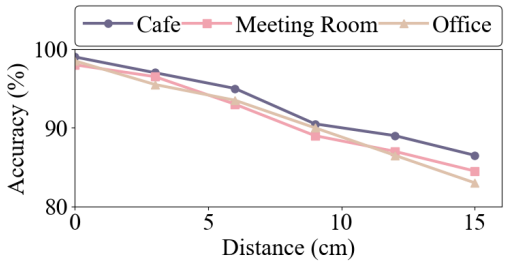

Fig. 7 The impact of attack distance on EMLogger.

6.2 Impact of Attack Distance

We use the disk inside the computer Lenovo Xiaoxin 14pro to evaluate the impact of distance between the attack device and the disk in three different environments. Particularly, the attack device is placed under the computer. We set the attack distance range from 0 cm to 15 cm in a step of 3 cm. The training set is collected at 0 cm, and we calculate the classification accuracy of activities at other attack distances. As shown in Fig 7, EMLogger achieves the highest accuracy when both the test set and training set are collected at the same position (i.e., 0 cm). As the attack distance increases, the accuracy decreases slightly. However, within the attack distance range of 15 cm, EMLogger can still guarantee a high accuracy (within 83%). These results demonstrate that the features extracted by EMLogger from EM signals are robust to distance variations.

7 Conclusion

In this paper, we propose a new EM side-channel attack targeting the disks inside the computers, namely EMLogger. We first build a theoretical model to show that the computer activities can be inferred by the EM signals leaked from the disks. To improve the SNR of EM signals, we propose the sub-clock fusion method. Then ,we design a DNN model to extract features from the EM signals and then achieve activity classification. Our real-world experiments show that EMLogger can effectively infer the computer activities from the EM signals and is robust to varying attack distances.

Acknowledgements. This paper is supported by the National Natural Science Foundation of China under grant U21A20462 and 62372400, “Pioneer” and “Leading Goose” R&D Program of Zhejiang under grant No. 2023C01033, and the Postdoctoral Fellowship Program of CPSF under Grant Number GZC20241488.

References

[1] Biedermann, S., Katzenbeisser, S., Szefer, J.: Hard drive side-channel attacks using smartphone magnetic field sensors. In: International Conference on Financial Cryptography and Data Security, pp. 489–496 (2015). Springer

[2] GOLD, B.D., LINDE, R.R., CUDNEY, P.F.: Kvm/370 in retrospect. In: 1984 IEEE Symposium on Security and Privacy, pp. 13–13 (1984). https://doi.org/10.1109/SP.1984.10002

[3] Quicke, S.: SSD becoming the norm in laptops. https://www.computerweekly.com/microscope/news/252478552/SSD-becoming-the-norm-in-laptops (2020)

[4] Morgan, D.: A Handbook for EMC Testing and Measurement vol. 8. Iet, ??? (1994)

[5] Alam, M., Khan, H.A., Dey, M., Sinha, N., Callan, R., Zajic, A., Prvulovic, M.: {One&Done}: A {Single-Decryption}{EM-Based} attack on {OpenSSL's}{Constant-Time} blinded {RSA}. In: 27th USENIX Security Symposium (USENIX Security 18), pp. 585–602 (2018)

[6] Genkin, D., Pachmanov, L., Pipman, I., Tromer, E., Yarom, Y.: Ecdsa key extraction from mobile devices via nonintrusive physical side channels. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1626–1638 (2016)

[7] Zhan, Z., Zhang, Z., Liang, S., Yao, F., Koutsoukos, X.: Graphics peeping unit: Exploiting em side-channel information of gpus to eavesdrop on your neighbors. In: 2022 IEEE Symposium on Security and Privacy (SP), pp. 1440–1457 (2022). IEEE

[8] Ni, T., Zhang, X., Zhao, Q.: Recovering fingerprints from in-display fingerprint sensors via electromagnetic side channel. In: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, pp. 253–267 (2023)

[9] Enev, M., Gupta, S., Kohno, T., Patel, S.N.: Televisions, video privacy, and powerline electromagnetic interference. In: Proceedings of the 18th ACM Conference on Computer and Communications Security, pp. 537–550 (2011)

[10] Chen, H.-W., Wu, J.-C.: A spread spectrum clock generator for emi reduction. IEICE transactions on electronics **84**(12), 1959–1966 (2001)

[11] Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. Communications of the ACM **60**(6), 84–90 (2017)

[12] Bjorck, N., Gomes, C.P., Selman, B., Weinberger, K.Q.: Understanding batch normalization. Advances in neural information processing systems **31** (2018)

[13] Hara, K., Saito, D., Shouno, H.: Analysis of function of rectified linear unit used in deep learning. In: 2015 International Joint Conference on Neural Networks (IJCNN), pp. 1–8 (2015). IEEE