

Replay-resistant Disk Fingerprinting via Unintentional Electromagnetic Emanations

Wenfan Song
Zhejiang University
Hangzhou, Zhejiang, China
wenfansong@zju.edu.cn

Yajie Liu
Zhejiang University
Hangzhou, Zhejiang, China
yajie@zju.edu.cn

Jianwei Liu
Zhejiang University
Hangzhou, Zhejiang, China
Hangzhou City University
Hangzhou, Zhejiang, China
jianweiliu@zju.edu.cn

Jinsong Han*
Zhejiang University
Hangzhou, Zhejiang, China
hanjinsong@zju.edu.cn

ABSTRACT

External disks (abbr., disks) are common data storage peripherals for hosts. Verifying the disk's legitimacy is crucial to prevent security issues on a host like privacy leakage and virus propagation before interaction setup. To address this issue, we propose *DiskPrint*, a novel non-intrusive and replay-resistant disk authentication system that relies on unintentional electromagnetic (EM) emanations from disks' internal components. The core idea of *DiskPrint* is that EM signals emitted during data writing can reflect hardware discrepancies among different disks. Based on electromagnetic principles, we establish a theoretical model associating EM signals with built-in electronic components to demonstrate the feasibility of extracting disk fingerprints from such EM emanations. We also propose a series of signal enhancement methods to remove the EM interface and improve the signal-to-noise ratio (SNR) of the EM measurements. To boost the security of *DiskPrint*, we propose a device-agnostic replay-resistant method by introducing randomness into leaked EM signals. Real-world experiments with 60 disks including hard disk drives (HDDs) and solid state drives (SSDs) from seven brands and 14 models indicate that *DiskPrint* achieves a 99%+ authentication success rate. Robustness analysis demonstrates *DiskPrint*'s stability over time. Security study shows its ability to defend against various attacks.

CCS CONCEPTS

• Security and privacy → Security in hardware.

KEYWORDS

Device Fingerprinting, Electromagnetic Radiation, Disk, Side-channel

*Jinsong Han is the corresponding author.

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. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

RAID 2024, September 30–October 02, 2024, Padua, Italy

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0959-3/24/09

<https://doi.org/10.1145/3678890.3678917>

ACM Reference Format:

Wenfan Song, Jianwei Liu, Yajie Liu, and Jinsong Han. 2024. Replay-resistant Disk Fingerprinting via Unintentional Electromagnetic Emanations. In *The 27th International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2024)*, September 30–October 02, 2024, Padua, Italy. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3678890.3678917>

1 INTRODUCTION

Nowadays, external disks (abbreviated as disks) are widely used to store digital content due to their large capacity, quick data transfer, and plug-and-play features. However, the interaction between a disk and a host usually does not require any authorization. This poses severe security problems against hosts. On the one hand, an adversary can directly use a disk to read (steal) data from a host (e.g., personal computer (PC)) [6], leading to privacy leakage. On the other hand, an adversary can leverage the disk to inject viruses [35] or malware into a host [20], once a user unwittingly plugs the infected disk into the computer. Even worse, the malicious code could propagate across the network, breaking more valuable systems [10].

The key to securing the host against the above threats is to control the access permissions of the disk. That is, before allowing “meaningful-data” interactions, the host needs to authenticate the identity and access level of the disk. In fact, people have already utilized a set of methods for device authentication, including the device identifier (e.g., serial numbers [51], vendor ID [17]), digital certificate [9], and password [50]. However, these methods are vulnerable to modification or mimicry attacks [31, 39, 44].

An effective countermeasure against the above limitations is to identify the device via its hardware fingerprint. The basis of hardware fingerprints is that the inherent physical characteristics of the devices are difficult to manipulate. However, existing hardware fingerprinting methods [2, 7, 11, 27, 55] cannot be trivially adopted for disk authentication. Specifically, these approaches can be divided into two categories. The first category relies on specific hardware modules (e.g., microphone [55], camera [27], or PUF [13, 43]) that disks do not initially possess. Such approaches require extra hardware or modifications on the disks [13, 27, 55], and thereby are inappropriate for the adoption in existing disks. The approaches in the second category are non-intrusive [11, 18, 41],

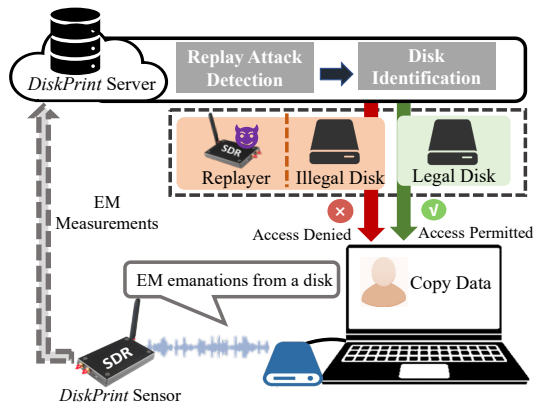


Figure 1: *DiskPrint* protects hosts from malicious disks.

requiring no extra hardware or modifications on the disks. They extract fingerprints from the devices' power consumption [41] and the wireless (e.g., radio frequency (RF) [11], electromagnetic (EM) [18]) signals reflected or leaked from the device. Nevertheless, these approaches exhibit various limitations: 1) Host-specific fingerprint [18, 41]. The fingerprints extracted by different models of hosts may differ, limiting their universality; 2) Vulnerability to replay attacks [11, 18]. The systems could be deceived by recorded-and-replayed signals from illegitimate devices; 3) Susceptibility to environmental interference [18]. They require 'clean' environments (e.g., an electromagnetic-safe zone) to collect quality fingerprints. These limitations hinder the practical deployment of these systems in practice.

In this paper, we propose *DiskPrint*, a novel non-intrusive and replay-resistant disk authentication system. It uses EM signals unintentionally leaked from disks to acquire and authenticate fingerprints. The key insight of *DiskPrint* is that, when writing data to a disk, its inherent hardware characteristics would be embodied by the EM signal it leaks. Such a signal varies among disks and thereby can serve as the disk fingerprint. Even if the disk is connected to different hosts, its fingerprints remain consistent, making them universally applicable across different hosts. By introducing randomness into the EM emanation, *DiskPrint* also bears the ability to resist replay attacks. As illustrated in Fig. 1, *DiskPrint* consists of an EM signal sensor, a stimulation program, and a server. It is noteworthy that a server is not necessary. *DiskPrint* can be directly integrated into the host without a server. Introducing a server allows all hosts to share the fingerprint database, improving availability and scalability. When a disk attempts to interact with a host, *DiskPrint* first interrogates its identity (ID). Then, a stimulation program triggers the disk to emit EM leakage, which is measured by the sensor and uploaded to the server. On the server side, *DiskPrint* first detects replay devices and then extracts disk fingerprints from received signals to verify the alleged disk ID. Ultimately, the host gives the disk corresponding permissions (e.g., acquiring data from the host) based on the verified ID.

Developing *DiskPrint* presents us with three challenges: (1) Generally, EM signal-based hardware fingerprint stems from the manufacturing imperfection of the device's inner electronic components.

A disk has numerous electronic components. It is hard to correlate the EM emanation and these electronic components, so as to prove that the leaked EM signals do contain sufficient hardware features for disk authentication. (2) Raw EM measurements have much noise that is irrelevant to the disk fingerprint, resulting in a low signal-to-noise ratio (SNR). It is difficult to extract quality and representative features from noisy and low-SNR EM signals. (3) In wireless systems, the prevention of replay attacks usually depends on randomizing the parameters of signals (e.g., amplitude [46]). Such signal-level randomization commonly requires full control over the device that generates the signal [23, 46]. However, in our scenario of disk authentication, the disk radiates signals unintentionally and we cannot perform signal-level modification. Hence, it is challenging to randomize the EM signals emitted by the disk to resist replay attacks.

We devise the following solutions to address the above challenges. (1) We analyze the working mechanism of the disk's internal modules during data writing, upon which we build a rigorous theoretical model associating the EM signal with the built-in electronic components based on electromagnetism. This model demonstrates that the leaked EM signals contain abundant physical characteristics of these electronic components, allowing us to extract qualified disk fingerprints. (2) To address the second challenge, we first propose a series of signal processing methods, including EM interference elimination and SNR enhancement, to enhance the signals. Then, we define several features to quantify the disk fingerprint and perform accurate disk classification. (3) To deal with the third challenge, instead of performing randomization at the signal level, we design a delicate application-level randomization mechanism. We find that manipulating the state of the disk's write caching (i.e., on/off) would lead to distinct variation scales in the leaked EM signals. By randomly selecting a one-time on-off sequence as the security credential and toggling the write caching accordingly, we can achieve the same defense effect as the signal-level one. That is, we can detect the replayed signal once the on-off sequence extracted from the real EM signal is different from the security credential.

To evaluate the performance of *DiskPrint*, we build a prototype with commercial off-the-shelf (COTS) devices and conduct real-world experiments using 60 disks with seven brands and 14 models. The experiment results indicate that *DiskPrint* can accurately identify legal disks with 99%+ success rates. Robustness analysis demonstrates that *DiskPrint*'s performance does not degrade over time. Security study manifests that *DiskPrint* is able to defend against various attacks.

In summary, our contributions are as follows: (1) We design and implement *DiskPrint*, to our best knowledge, the first non-intrusive and replay-resistant disk authentication system. *DiskPrint* can also be used to fingerprint internal disks. (2) We build a rigorous theoretical model to demonstrate that the EM emanations of a disk contain abundant physical features of its inner electronic components. We believe that the insight behind *DiskPrint* would be helpful in mining distinguishable features for fingerprinting other potential devices. (3) We propose several signal processing techniques to obtain a clean fingerprint signal. Meanwhile, we also design an effective disk classification method to achieve accurate disk authentication. (4) We evaluate the performance of *DiskPrint* on 60 real disks covering mainstream brands and models. The results show that *DiskPrint*

achieves over 99% authentication success rate. Meanwhile, *DiskPrint* is secure against various attacks.

2 THREAT MODEL

In the considered attack scenario, the attacker uses an unauthorized disk to steal or damage the host’s data through virus injection and etc. To protect the host, we introduce *DiskPrint* for disk authentication before interaction. However, if the attacker knows *DiskPrint*’s deployment, s/he may mount more adaptive attacks like mimic and replay attacks to make *DiskPrint* accept unauthorized disks, and compromise the host.

Mimic Attack. We assume that the attacker is able to acquire a substitute disk (Disk B) that is similar to the target disk (Disk A), e.g., they belong to the same brand and model. The attacker tries to spoof *DiskPrint* by initiating an authentication request with B, expecting that *DiskPrint* mistakes Disk B for A and grants Disk B the permissions of A.

Replay Attack. We assume that the attacker has necessary attack devices, including a device to receive the EM signal with the same frequency band as that of the disk emanation and a replay device. The attacker first records the EM signal during a legal authentication attempt (of Disk A). Then, the attacker launches an illegal authentication request with an illegal disk (Disk B). During Disk B’s EM emanation generation, the attacker blocks it with a metal casing but replays the recorded EM signal to *DiskPrint*’s sensor, aiming to trick *DiskPrint* into granting Disk B the permissions of Disk A.

We also assume that the communication between the EM sensor and the server and between the server and the host is secure, which can be achieved by secure communication protocols and encryption techniques [29]. Besides, we do not consider the following cases: 1) The attacker deliberately breaks legitimate authentication, such as jamming EM sensors and physically destroying legal disks. 2) A person uses legal disks to perform malicious activities, e.g., injecting malware into or physically modifying legal disks. In fact, *DiskPrint* can be used as a traceability tool to trace the malicious activities of legitimate disks by correlating the disk ID with such illegal operations.

3 FINGERPRINT IN EM EMANATION OF DISK

This section validates the feasibility of EM fingerprint by establishing a theoretical model and conducting preliminary experiments.

3.1 Theoretical Model

When data is written to a disk, the high activity of its inside electronic components would incur heavy currents and further induce EM signals. The clock module and dynamic random access memory (DRAM) module are the main sources of these EM signals, and the electrical characteristics of the electronic components in these modules predominantly determine the strength of the signals [5]. Due to the manufacturing imperfection, these electronic components contain “unique” features. As a result, the EM leaked from each disk will be unique as well, and can be used for fingerprinting. As the number of these electronic components is very large [49], the extracted fingerprint will be highly distinguishable. We theoretically detailed the basis of this feature exhibition as follows.




Figure 2: Structure of the clock module and the DRAM module in a disk.

We assume that an I/O operation of a disk, such as writing data, induces EM signals $s_{data}(t)$. It is primarily composed of two parts: one from the clock module ($s_{clk}(t)$) and the other from the DRAM module ($s_{dram}(t)$). As the write operations of the DRAM module are performed on the clock edges, $s_{dram}(t)$ will be amplitude-modulated to $s_{clk}(t)$, leading to the eventually leaked signal $s_{data}(t)$:

$$s_{data}(t) = s_{dram}(t)s_{clk}(t). \quad (1)$$

We will explore the fingerprint in $s_{data}(t)$ by analyzing the signals leaked by the clock and DRAM modules, respectively.

■ Fingerprint From Clock Module. The clock module is an essential component used to synchronize the operation of various components in a disk. To comply with the electromagnetic compatibility (EMC) regulations [8], the spread spectrum clocking (SSC) technology is widely used in current disks [42]. As shown in Fig. 2, the clock module typically consists of four parts: an oscillator circuit (OSC), a phase-locked loop (PLL), a frequency modulator (FM), and a band-pass filter (BPF). The signal $s_{clk}(t)$ leaked from the SSC clock is a sum of signals distributed at frequency $f_0 - if_m$ ($0 \leq i < k$), which can be expressed as:

$$s_{clk}(t) = \sum_{i=0}^{k-1} A_i \cos(2\pi(f_0 - if_m)t). \quad (2)$$

In this equation, A_i is the amplitude of the i -th sub-clock and is determined by the band-pass filter and frequency modulator. Due to the hardware imperfections created during the manufacturing process, there are differences in these components among different clock modules, leading $s_{clk}(t)$ of different disks to be distinguishable.

■ Fingerprint From DRAM Module. Within DRAM modules, millions of complementary metal oxide semiconductor (CMOS) transistors are arranged in a lattice formation to constitute DRAM banks for temporal data storage [34]. A unit of the DRAM bank is called a cell, which consists of a CMOS transistor and a capacitor, as shown in Fig. 2. During the data writing process, the continuous charging and discharging of CMOS transistors in cells generate heavy currents. According to Maxwell’s equations [4], the varying currents will create an EM field that accompanies the EM signal $s_{dram}(t)$. That is, $s_{dram}(t)$ is primarily determined by the varying currents that are directly related to the dynamic power consumption $P_{dynamic}$ of the CMOS circuits [28]. Particularly, $P_{dynamic}$ can be expressed as:

$$P_{dynamic} = \frac{1}{2}CV_{DD}^2f_0\alpha, \quad (3)$$

where C is the equivalent capacitance that is determined by the CMOS transistors. V_{DD} denotes the supply voltage which is determined by the voltage regulator and power controller in the DRAM

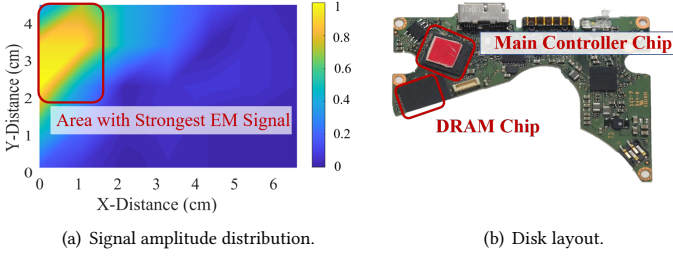


Figure 3: Clock module and DRAM module radiate strongest EM signals (the clock module is in the main controller chip).

module, f_0 represents the frequency of the clock signal that is determined by the clock module, and α refers to the activity factor that is determined by the workload of the DRAM module. During the fabrication, the hardware imperfection of relevant electronic components in DRAM would induce the uniqueness in V_{DD} , f_0 , and C , and hence radiate distinguishable s_{dram} .

In summary, since s_{clk} and s_{dram} can embody the unique hardware characteristics of disks, s_{data} has the potential to serve as the disk fingerprint.

3.2 Preliminary Experiment

We conduct several preliminary experiments to further verify that (1) the EM signal $s_{data}(t)$ leaked by the disk mainly originates from the clock module and DRAM module, (2) $s_{data}(t)$ from different disks is distinguishable, and (3) $s_{data}(t)$ is stable enough to serve as a qualified fingerprint carrier. In the experiments, we employ a stimulation program to induce the disk to emit EM signals, which are collected by a software-defined radio (HackRF [30]), a Foresight low-noise amplifier (FSTRFAMPO1 [19]), and a 3dbi antenna [1]. Each collection lasts 1.5 seconds. One issue here is that the workload of the DRAM module also affects $s_{data}(t)$, as mentioned in Sec. 3.1. To avoid the impact of diverse workload, the stimulation program adopts a uniform workload, i.e., repeatedly writing data to a fixed area in the disk at its maximum speed.

■ **EM Radiation Source Validation.** To confirm that the EM signal leaked by the disk is mainly from the clock module and DRAM module, we measure the signal strength by alternatively placing the antenna on 70 spots of a Western Digital (WD) Elements SE disk’s surface. Figure 3(a) presents the heat map of the signal strength. Combined with the layout of the disk’s printed circuit board (PCB) shown in Fig.3(b), we can conclude that the area with the strongest EM signal in the heat map corresponds to the clock module and DRAM module. This demonstrates that the leaked EM signals are indeed primarily from these two modules.

■ **Evidence of Disk Fingerprint.** To validate that the discrepancy between different disks can be revealed by the EM signals they leak, we collect the EM signals (i.e., $s_{data}(t)$) of four disks, including two WD Elements SE (WD Elements SE-1 and WD Elements SE-2), a TOSHIBA DTB410, and a WD My Passport Ultra. As $s_{data}(t)$ includes multiple sub-clock signals that are similar to each other, we only show one (the i -th) normalized sub-clock signal (i.e., $|A_i s_{dram}(t) \cos(2\pi(f_0 - if_m)t)|_{normalized}$) in Fig. 4. We observe




Figure 4: EM signals of four disks. Nor. Amp. means the normalized amplitude.

that there are obvious amplitude profile differences¹ between any two disks with different models (Fig. 4(a) vs. (c)/(d)). Meanwhile, the sub-clock signals of two disks with the same model ((Fig. 4(a) vs. (b))) also demonstrate noticeable discrepancies. This confirms the feasibility of using the EM leakage for disk authentication.

■ **Temporal and Spatial Consistency.** A qualified fingerprint should be temporally and spatially stable. To investigate the temporal and spatial consistency, we collect EM signals of one disk (i.e., WD Elements SE-1) at different locations (laboratory and office) on different days (Day 1 and Day 2 with one week in between). The normalized signals collected on Day 1 in the laboratory, Day 2 in the laboratory, and Day 1 in the office are shown in Fig. 5(a), (b), and (c), respectively. It can be observed that the amplitude profiles of sub-clock signals collected on different days are highly similar (Fig. 5(a) vs. (b)). Also, the amplitude profile of the sub-clock signal does not exhibit significant changes when the location varies (Fig. 5(a) vs. (c)).

These experiments demonstrate the feasibility of our theoretical analysis. The amplitude profiles of the EM signals leaked by disks have high distinguishability and reliability, making them competent to serve as the disk fingerprint carrier.


4 DISKPRINT DESIGN

4.1 System Overview

Figure. 6 illustrates the workflow of *DiskPrint*, which consists of two main phases: registration and authentication.

■ **Registration Phase.** To register a new disk, *DiskPrint* first utilizes a stimulation program to make it generate EM signals containing fingerprint information (called fingerprint signals) to obtain registration data. Then, on the server, *DiskPrint* enhances the signals by removing EM interference via convolutional filtering and improving the SNR using a clock jitter-resistant folding algorithm (detailed in Sec. 4.2). Thereafter, *DiskPrint* conducts a series of pre-processing and extracts feature vectors quantifying the fingerprint from the registration data to form a positive training set. With the negative training set (of other disks) provided by *DiskPrint*’s server, *DiskPrint* trains a binary disk classifier (named disk identifier) for

¹We consider amplitude- instead of frequency-exhibited characteristics as the $s_{data}(t)$ ’s frequency distributions of the same model of disks are similar and frequency jitters sometimes, which may cause misclassification.



(a) In laboratory on Day 1. (b) In laboratory on Day 2. (c) In office on Day 1.

Figure 5: EM signals of one disk collected in different places (laboratory and office) on different days (Day 1 and Day 2 with one week in between) are consistent. Nor. Amp. means the normalized amplitude.

this disk. The disk identifier and disk ID (e.g., serial number) will be stored in the server (detailed in Sec. 4.3).

■ **Authentication Phase.** When a disk attempts to interact with a host, *DiskPrint* first interrogates the disk’s ID. Then, *DiskPrint* uses the stimulation program to generate the EM trace. This trace is composed of two parts: the signal segment used for replay detection (termed as preamble signal) and the fingerprint signal. After performing the signal enhancement, *DiskPrint* determines if the received signal is from a replay device by checking the preamble signal (detailed in Sec. 4.4). If not, the fingerprint signal will go through the same pre-processing and feature extraction as the registration phase. Thereafter, the extracted feature vector will be fed into the pre-trained disk identifier associated with its ID. If the output of the disk identifier is negative, the fingerprint signal is deemed to be from an illegal disk and this disk will be rejected. Otherwise, *DiskPrint* confirms the authenticity of the ID, and this disk will be authorized to interact with the host.

4.2 Signal Enhancement

As mentioned in Sec. 3.2, the disk fingerprint is derived from the amplitude profiles of the sub-clock signals. However, this process is susceptible to interference from other EM sources in ambient environments. Besides, each disk should pass the EMC test according to the regulation [8]. Yet, following such requirements lead the strength of the raw fingerprint signal to be relatively low. Compared to ambient interference, the weak signal strength results in a low SNR. Both of the above factors would degrade the quality of extracted fingerprints and reduce authentication accuracy. To solve these problems, we propose two signal enhancement techniques: EM interference elimination and clock jitter-resistant folding. Note that we perform EM interference elimination before clock jitter-resistant folding in *DiskPrint*’s workflow. But in the following, we first introduce the folding to better clarify the effectiveness of the EM interference elimination.

4.2.1 Clock Jitter-resistant Folding. To address the low SNR of fingerprint signals, we focus on the correlations among sub-clocks. For each disk, we noticed two phenomena: (1) the frequency interval between adjacent sub-clocks is fixed (i.e. f_m) and (2) the amplitude profiles of different sub-clocks are consistent. These phenomena inspire us to use the folding algorithm [26] to improve the fingerprint signal’s SNR by accumulating the sub-clocks on the spectrum. Nevertheless, in practice, the frequencies of the sub-clocks jitter within a small range, although they should remain unchanged in

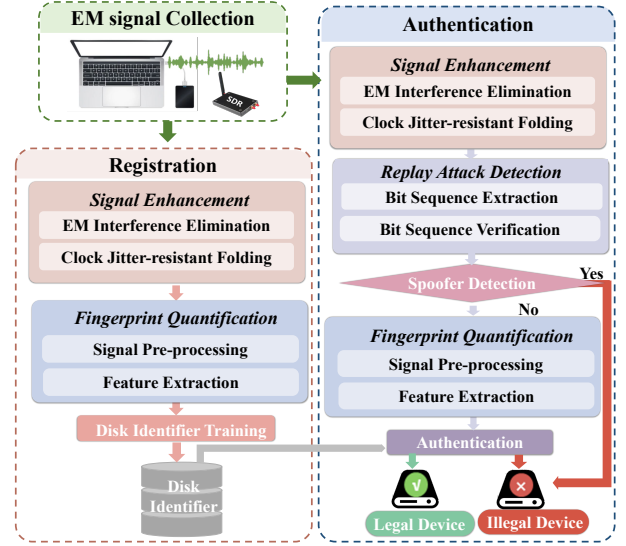


Figure 6: Workflow of *DiskPrint*.

theory. In the following, based on the original folding algorithm, we propose a countermeasure to achieve clock jitter-resistant folding.

Specifically, we first perform a short-time Fourier transform (STFT) to obtain the spectrogram X of the fingerprint signal $s_{data}(t)$. The STFT applies a Hanning window w to divide $s_{data}(t)$ into small segments with the same length, and performs a N -point discrete Fourier transform (DFT) on each segment. Suppose that the sampling rate is f_s , then the frequencies of adjacent sub-clocks will be separated by $\theta = f_m \times \frac{N}{f_s}$ DFT bins, and the clock frequency f_0 locates in the h -th DFT bin. The STFT’s result can be expressed as:

$$\text{STFT}[s_{data}(t)] = \text{DFT}[w] * \sum_{i=0}^{k-1} s_{dram}(t) A_i \delta[h - i \times \theta], \quad (4)$$

where δ represents the impulse function and $s_{dram}(t) A_i$ reflects the amplitude profile of the i -th sub-clock. Then, the spectrogram X of the fingerprint signal can be expressed as:

$$X = |\text{STFT}[s_{data}(t)]|^2, \quad (5)$$

Let $P[j, t]$ denote the folding result at time t , i.e., the sum of the amplitude of the j -th, $(j - \theta)$ -th, \dots , $[j - (k - 1) \times \theta]$ -th DFT bins in the spectrogram X :

$$P[j, t] = \sum_{i=0}^{k-1} X[j - i \times \theta, t]. \quad (6)$$

If the clock frequency f_0 is located in the h -th DFT bin, the strength of $P[h, t]$ will be much higher than that of $P[j, t]$ when $j \neq h$. This is because the correlations between sub-clocks would strengthen the signal power ($P[h, t]$), while noise ($P[j, t]$, $j \neq h$) will be overwhelmed owing to its randomness. Thus, we can effectively improve the SNR of the fingerprint signal through folding.

In reality, the frequencies of sub-clocks will jitter over a small range (the upper limit is f_m), rendering it difficult to locate sub-clocks. Fortunately, the jitters of all the sub-clocks are synchronous.




Figure 7: Effectiveness of EM interference elimination. Nor. Amp. means normalized amplitude.

This makes the interval f_m between sub-clocks remain constant. Besides, the folding result of sub-clocks is much higher than that of the noise. With these properties, to locate the frequencies of the sub-clock signals, we can find the maximum value of $p[j, t]$ in the range of $[h - 0.5\theta, h + 0.5\theta]$, and update h to the one whose folding result is significantly larger than others.

4.2.2 EM Interference Elimination. As aforementioned, the folding algorithm can enhance the SNR of the fingerprint signals. But the EM interference in ambient environments may submerge the fingerprint signal, making it difficult to locate the sub-clock signals in the spectrogram and disabling the folding algorithm. To tackle this problem, we propose a convolution filter-based denoising method to remove the noise from the spectrogram. This effectively mitigates the impact of EM interference on the folding algorithm.

Particularly, the actually collected EM signals consist of both the fingerprint signal and EM noise. The corresponding spectrogram $X'[j, t]$ can be expressed as:

$$X'[j, t] = X[j, t] + N[j, t] \quad (7)$$

where $X[j, t]$ denotes the spectrogram of the “clean” fingerprint signal and $N[j, t]$ is the spectrogram of noise. To remove $N[j, t]$ from $X'[j, t]$, we propose to convolve the spectrogram $X'[j, t]$ with a kernel of $[-0.5, 1, -0.5]$. In this way, we can get a new spectrogram:

$$\tilde{X}[j, t] = -\frac{1}{2}X'[j-1, t] + X'[j, t] - \frac{1}{2}X'[j+1, t]. \quad (8)$$

Assuming that the clock frequency f_0 is located in the h -th DFT bin of the spectrogram $\tilde{X}[j, t]$, then, the folding result $P'[h, t]$ becomes:

$$P'[h, t] = \sum_{i=0}^{k-1} \tilde{X}[h-i \times \theta, t] = P[h, t] + Q[h, t], \quad (9)$$

where

$$Q[h, t] = \sum_{i=0}^{k-1} \{N[h-i \times \theta, t] - \frac{1}{2}N[h-i \times \theta - 1, t] - \frac{1}{2}N[h-i \times \theta + 1, t]\}.$$

Since the element in N (i.e., $N[j, t]$) is independent of each other, i.e., N is a random sequence, it is evident that $Q[h, t]$ is statistically close to zero and much smaller than $P[h, t]$ (i.e., the folding result of the fingerprint signal). Hence, we have $P'[h, t] \approx P[h, t]$, which indicates that the EM interference is effectively suppressed.

To validate the efficacy of our proposed EM interference elimination method, we conduct a validation experiment. Specifically, we collect the fingerprint signals of a Toshiba DTB410 disk in both clean and EM interference-rich environments, referred to as “clean signal” and “noisy signal”, respectively. Then, we utilize the EM interference elimination method to process the noisy signal and generate a “denoised signal”. The spectra of the noisy, denoised, and clean signals are shown in Fig. 7(a), (b), and (c), respectively. It can be found that the EM interference makes it difficult to locate the sub-clock signals in the spectrum (Fig. 7(a)), whereas the sub-clock signals in the denoised signal (Fig. 7(b)) and clean signal (Fig. 7(c)) can be precisely located. Additionally, we also show the folding results of the noisy signal, denoised signal, and clean signal in Fig. 7(d), (e), and (f), respectively. Due to the strong EM interference, the folding result of the noisy signal (Fig. 7(d)) is severely distorted, but the folding result of the denoised signal (Fig. 7(e)) is similar to that of the clean signal (Fig. 7(f)). These results prove that our EM interference elimination method can effectively suppress the impact of EM interference on the folding result.

4.3 Feature Extraction and Disk Classification

This subsection describes how *DiskPrint* extracts features from the folding result $P'[h, t]$ (simplified as $P'[t]$) and performs disk classification.

■ **Signal Pre-processing.** Due to instantaneous interference in the circuit, etc., the current in the circuit could change sharply, resulting in anomalies with extremely large/small values in $P'[t]$. To eliminate their impacts, we leverage a medium absolute deviation (MAD)-based outlier detection method [48] to find out all anomalies, and then remove them from $P'[t]$. Moreover, considering that the position of *DiskPrint*'s sensor with respect to the disk could change in different authentication attempts and the gain of different sensors

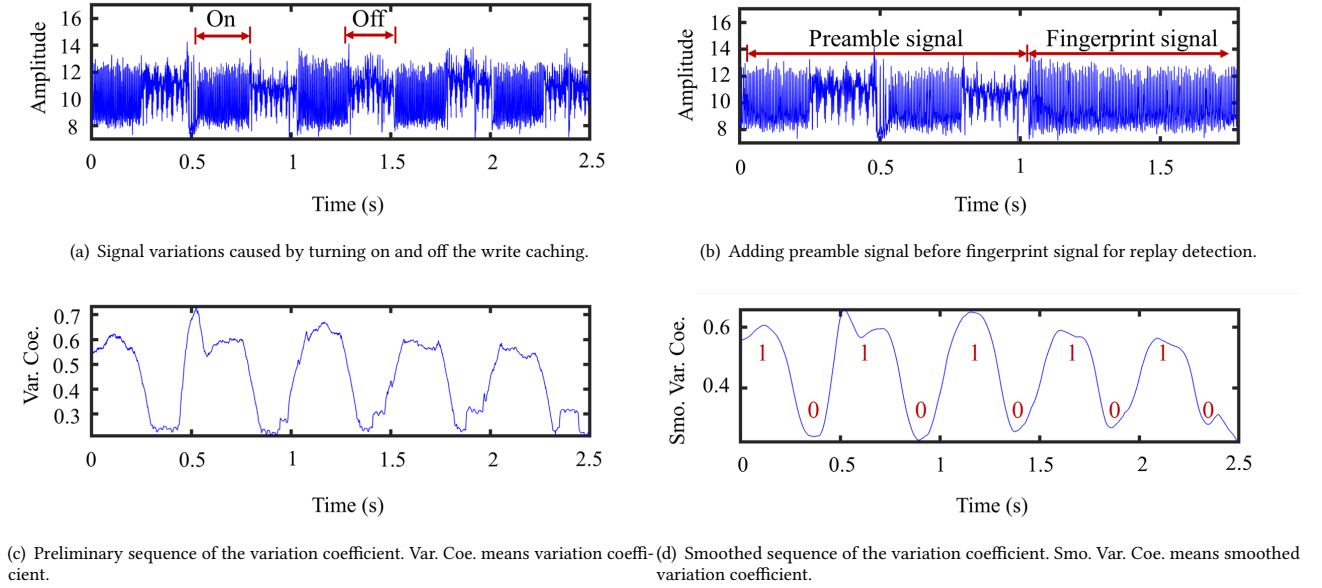


Figure 8: Anti-replay method.

could vary, the strength of two folding results $P'[t]$ of the same disk may differ. To deal with this problem, we adopt min-max normalization to limit the range of $P'[t]$ to $[0, 1]$. Each normalized element in $P'[t]$ can be calculated by:

$$P'_{norm}[t] = \frac{P'[t] - P'_{min}}{P'_{max} - P'_{min}}, \quad (10)$$

where P'_{min} and P'_{max} are the minimum value and the maximum value of $P'[t]$, respectively.

■ **Feature Extraction.** According to Sec. 3.2, we observe significant differences in the amplitude profiles of EM signals leaked by different models of disks, even those leaked by the same model. Therefore, we select twelve features [32] that can describe the amplitude profile of the signal as the disk fingerprint, including *mean value*, *median value*, *standard value*, *variation coefficient*, *median absolute deviation*, *mean absolute deviation*, *root mean square*, *skewness*, *kurtosis*, *form factor*, *crest factor*, and *pulse factor*. To comprehensively mine the uniqueness of leaked signals, we extract such twelve features from both the time series and spectrum of $P'_{norm}[t]$. These 24 features form a feature vector and will be used as the input of the disk identifier.

■ **Disk Classification.** *DiskPrint* adopts supervised learning to achieve disk classification. Specifically, each registered disk will provide k feature vectors as positive samples. Then, the server of *DiskPrint* would also provide k feature vectors from other registered disks as negative samples. These samples form a training set for the disk identifier. We employ an ensemble learning method, random forest (RF), as the disk classifier owing to its high robustness, high speed, and ability of understanding feature importance [36]. In this way, each disk would have its own binary classifier (which only needs to be trained once) associated with its ID.

When a disk attempts to interact with the host, it should show its ID to *DiskPrint* at first. Then, *DiskPrint* feeds the feature vector into the classifier associated with this ID to verify the authenticity of the disk's purported identity.

4.4 Replay Detection

Replay attack is one of the most intractable threats against EM-based sensing/communication systems [47]. Traditional methods to prevent replay attacks typically rely on randomizing the parameters of signals (e.g., amplitude [46]), which usually requires complete control over the signal generation equipment [23, 46]. However, in our disk authentication scenarios, we are unable to directly perform signal-level modifications to EM signals unintentionally radiated by the disk. To tackle this problem, this subsection proposes a device-agnostic replay-resistant method by introducing application-level randomness into the disk's EM emanations.

■ **Insight.** We observed that the write caching of a disk can be in either the enabled (on) or disabled (off) state during data writing. The state switching of the write caching would change the access frequency of the DRAM module and consequently affects its workload [37]. As a large/small workload would induce large/small power of the DRAM module, different states produce EM traces with distinct variation scales (shown in Fig. 8(a)). Based on this phenomenon, we can employ the stimulation program to manipulate the state of the disk's write caching at the application level and introduce controllable randomness into the EM leakage to detect replay attacks.

■ **Replay-resistant Design.** We first design a binary encoding scheme for the EM trace used for replay detection. Turning on/off the write caching can generate EM signals with large/small discreteness that can be encoded as 1/0, or vice versa. By alternately turning on and off the write caching within a period of time, *DiskPrint* can




Figure 9: Experiment Setup for performance evaluation.

obtain a “10111...” like bit sequence. Next, we explain how to identify replayed signals through bit sequence verification. Prior to generating fingerprint signals, *DiskPrint* first randomly selects a one-time bit sequence as the security credential. Meanwhile, we reserve a period for the stimulation program to embed the security credential into the EM trace by accordingly turning on or off the write caching. The corresponding EM signal is the called preamble signal (see Fig. 8(b)). After receiving the preamble signal, *DiskPrint* extracts a bit sequence based on the encoding scheme. If the extracted bit sequence differs from the security credential, the authentication request is very likely initiated by a replay attacker. In our default setting, a bit sequence consists of 10 bits, giving an encoding space of $2^{10} = 1024$. Theoretically, once the extracted bit sequence is not the same as the security credential, we are $1 - \frac{1}{1024} = 99.9\%$ sure that this is a replay attack.

■ **Bit Sequence Extraction.** It can be found that the one-time and random bit sequence is the key enabler of our replay-resistant method. Inaccurate extraction of bit sequence would cause false rejects of legal authentications. Therefore, we develop an effective method to accurately extract the bit sequence from noisy preamble signals. To be specific, *DiskPrint* first removes discrete points in the preamble signals, which can be regarded as sharp noise. Then, to eliminate the impact of the position of *DiskPrint*’s sensor with respect to the disk on the signal strength, *DiskPrint* normalizes the preamble signal. Next, to determine whether a segment in the preamble signal belongs to “0” or “1”, we leverage the variation coefficient (defined in Sec. 4.3) to quantify the degree of discreteness. In particular, we identify each bit by detecting the boundary between high and low variation coefficients in the following steps: (1) Sliding a window on the normalized preamble signal and calculating the variation coefficient of each window. This yields a preliminary sequence of the variation coefficient D , as shown in Fig. 8(c). (2) Smoothing the sequence D by mean-filtering and get a smoothed sequence D' , as shown in Fig. 8(d). (3) Setting an empirical threshold D'_{thr} to detect the boundary between high and low variation coefficients in D' . If an element with index i in D' satisfies $D'(i-1) < D'_{thr}$ and $D'(i+1) > D'_{thr}$, then this element is the boundary from “0” to “1”; otherwise, it is the boundary from “1” to “0”. With this approach, we can digitize the sequence D' and obtain a reliable bit sequence for replay detection.

Table 1: Disks utilized in the evaluation of *DiskPrint*.

Number	Disk Manufacturer	Model
1-15	Western Digital	WDBEPK0020BBK-EB
16-20	Western Digital	WDBU6Y0040BBK-EB
21-24	Western Digital	WDBFTM0050BGD-OD
25-27	Western Digital	WDBYVG0020BBK-0A
28	Western Digital	WDBAGF5000ASL
29	Western Digital	WDBAYN4800ABK
30-44	Toshiba	DTB410
45-48	Toshiba	DTB420
49-53	Seagate	SRD00F1
54	Seagate	SRD0FV4
55	SanDisk	SDSSDE30-480G
56	Aigo	S7
57-58	Lenovo	F308
59-60	Newsmy	Qingfeng

5 EVALUATION

In this section, we present the implementation and evaluate the overall performance, robustness, and security of *DiskPrint*.

5.1 Experiment Setup

■ **Hardware.** As shown in Fig. 9, we utilize a software-defined radio (HackRF SDR), a Foresight low noise amplifier (FST-RFAMP02), and a 3dbi antenna to constitute the sensor of *DiskPrint*. The amplifier is connected to HackRF, while the antenna is connected to the amplifier and placed on the surface of the disk being tested. We employ a HUAWEI KLV-W19L PC with Intel(R) Core(TM) i5-8265U CPU running Ubuntu 18.04 OS as the default host for disk interaction. This PC is also used to process the received signals.

■ **Stimulation Program.** When inducing the disk to generate preamble and fingerprint signals, the stimulation program repeatedly writes data “0” to a fixed area in the disk at the maximum speed of this disk. As described in Sec. 4.4, the stimulation program encodes the preamble signal by alternatively switching the state of the write caching. When generating fingerprints, the write caching is maintained in the “on” state. In consideration of the universality and compatibility, we use C++ language to implement the stimulation program.

■ **Data Collection.** We conduct experiments in three different environments: laboratory, office, and home. For each environment, we evaluate *DiskPrint* on 60 disks (shown in Table. 1), including both HDDs and SSDs with 7 popular brands [45] (Western Digital, Toshiba, Seagate, Newsmy, Aigo, Lenovo, and SanDisk). At most 15 disks belong to the same model. The sampling rate is set to 20 MHz. Considering that the leaked EM signals are typically distributed within the frequency range of 10 MHz and above, we set the bandwidth to 10 MHz. We collect over 65500 fingerprint signals, with each disk providing at least 900 fingerprint signals. Each fingerprint signal can correspond to an authentication attempt, which contains 1000 sampling points collected within 1.64 seconds.




Figure 10: ASRs of *DiskPrint* and two baselines.




Figure 11: FARs of *DiskPrint* and two baselines.




Figure 12: Performance of five classifiers.




Figure 13: Effect of the length of fingerprint signal.




Figure 14: Effect of training set size.




Figure 15: Effect of Host.

Metric. Four metrics [52] are defined to quantify the performance of *DiskPrint*: authentication success rate (ASR), false accept rate (FAR), defense success rate (DSR), and false reject rate (FRR). ASR represents the probability that *DiskPrint* successfully accepts a legal disk. FAR is the probability that *DiskPrint* falsely accepts an illegal disk. DSR means the probability that *DiskPrint* successfully detects an attack and FRR represents the probability that *DiskPrint* mistakenly rejects a legal disk. Higher ASR and lower FRR indicate that *DiskPrint* has better usability and user-friendliness, while lower FAR and higher DSR mean better security.

5.2 Overall Performance

In this part, we first evaluate the effectiveness of our signal enhancement methods, and then investigate the effect of the training set size, the scalability of *DiskPrint* and the effect of alien devices. Unless specifically stated, otherwise, for each disk, we randomly select 60 positive samples and 60 negative samples (from the other 59 disks' fingerprint signals) to construct the training set, while the rest fingerprint signals form the test set. For each disk identifier, there is no sample overlap between the training and test sets. The final result is taken as the mean of ten times of classification trials.

■ **Effectiveness of Signal Enhancement.** To evaluate the effectiveness of the two techniques used for signal enhancement, i.e., the EM interference elimination and clock jitter-resistant folding, we compare *DiskPrint* with two baselines. Baseline 1: the raw EM signals are directly collected by *DiskPrint*'s sensor (without the signal enhancement). Baseline 2: the signals are processed only through the clock jitter-resistant folding. Fig. 10 and Fig. 11 show ASRs and FARs for *DiskPrint* and the two baselines, respectively. The results of Baseline 1 indicate that the average ASR is below 86% and the average FAR is above 7% when directly using raw signals for authentication. The relatively low ASR and high FAR are due to the presence of electronic devices (such as computers), which introduce electromagnetic interference into the raw signals collected in these three environments. With clock jitter-resistant

folding (i.e., Baseline 2), the ASRs increase to 93.0%, 92.0%, and 93.2% for the three environments, while the FARs decrease to 2.4%, 2.5%, and 2.5%, respectively. Compared such results with those of Baseline 1, we can find that our clock jitter-resistant folding method is effective in improving the signal quality. Besides, adopting both of the two signal enhancement methods increases the average ASR to 99.5% and decreases the average FAR to 1.5%. This means that our EM interference elimination method also performs very well in suppressing electromagnetic interference in the surrounding environment as well as enhancing the signal. Moreover, we can notice that both the ASRs and FARs of the three environments are similar to each other. Thus, *DiskPrint* can accurately authenticate disks under different environments.

■ **Performance of Different Classifiers.** To demonstrate the superiority of our selected disk identifier, namely the RF classifier, we compare it with four commonly-used classifiers [22]: support vector machine (SVM), naive Bayes (NB), K-nearest neighbours (KNN), and decision tree (DT). The ASRs and FARs are shown in Fig. 12. It can be observed that NB has the lowest ASR of 84.3%, while SVM shows the largest FAR of 32.6%. KNN and DT demonstrate high ASRs, i.e., 95%+, but their FARs are high as well. Our disk identifier owns the highest ASR of 99.5%. Meanwhile, it performs very well in rejecting illegal disks with a low FAR of 1.5%. These results demonstrate that RF outperforms other classifiers.

■ **Effect of Fingerprint Signal Length.** In our default setting, one fingerprint signal contains 1000 sampling points, i.e., the length of each fingerprint signal is 1000. To evaluate the appropriate length, we vary the number of sampling points from 100 to 1000 and calculate the corresponding ASRs and FARs. The results shown in Fig. 13 indicate that the ASR increases with the fingerprint signal length, while the FAR is the opposite. This is reasonable as, intuitively, more sampling points can carry more stable and plentiful features of the disk. However, it is noteworthy that even with 100 sampling points (costs about 0.164 seconds), *DiskPrint* can already achieve 91.0% ASR and 9.0% FAR. When the fingerprint signal length increases to




Figure 16: Cross-occupation rate performance.




Figure 17: Cross-background application performance.




Figure 18: Cross-position performance.

800, the ASR and FAR can respectively reach 99.3% and 2.0%, which approaches the performance under 1000 sampling points. Since sampling 1000 points only takes about 1.64 seconds and provides the best performance, we recommend it as the default choice.

■ **Effect of Training Set Size.** In this experiment, for each disk, we vary the number of positive fingerprint signals from 10 to 60 in step of 10 (with an equal number of negative samples). We next recalculate the ASRs and FARs. Fig. 14 demonstrates that the ASR/FAR increases/decreases with the training set size. When only 10 training samples per disk are used, the ASR can already reach 94.4%. When the training set size increases to 50, the ASR/FAR increases/decreases to 99.0%/2.0%. This indicates that *DiskPrint* can accurately identify disks with only a small dataset. The highest/lowest ASR/FAR, 99.7%/1.6%, can be achieved when we use 60 training samples per disk. It is worth noting that *DiskPrint* takes less than two minutes to collect 60 fingerprint signals, highlighting its outstanding user-friendliness.

■ **Scalability Assessment.** To explore if *DiskPrint* can authenticate a large number of devices, similar to the experimental methodology in [7], we observe the variation trend of the performance with the number of tested devices. For doing so, we vary the number of involved disks (i.e. participating in training and test) from 15 to 60 in the step of 15 and recalculate the ASRs and FARs. The resulting ASRs for 15, 30, 45, and 60 disks are 99.6%, 99.6%, 99.5%, and 99.5%, respectively. Meanwhile, the FARs for these disk numbers are 1.3%, 1.4%, 1.5%, and 1.5%, respectively. It is apparent that the performance of *DiskPrint* remains stable with the number of tested disks. This is reasonable as the disk fingerprint has high distinguishability (Sec. 3.1). The result is an encouraging sign that *DiskPrint* can support a larger disk set.

■ **Effect of Alien Device.** In reality, *DiskPrint* may need to authenticate alien disks, i.e., unregistered disks. To understand the performance of *DiskPrint* towards alien disks, we randomly select 30 disks to register and the other 30 disks as unregistered disks. Then, the 30 alien disks take turns to input their fingerprint (50 instances per disk) into each of the classifiers of registered disks. The FAR of these $30 \times 30 \times 50 = 45000$ times of tests, 1.7%–, indicates the high reliability of *DiskPrint*.

5.3 Robustness Analysis

■ **Effect of hosts.** We also investigate the universality of fingerprints across different hosts: whether the fingerprint remains consistent when the disk is connected to different hosts. To explore the impact of hosts on *DiskPrint*'s performance, we collect fingerprint signals from four different models of hosts (HUAWEI KLV-W19L,

DELL Latitude E5570, ASUS A555L, and DELL OptiPlex 7060). Given that a host interacting with the disk may install any compatible operating systems (OSes), we simultaneously explore the influence of OS on *DiskPrint*. Accordingly, we deploy one of the following OS versions on each of the four hosts: Linux (Ubuntu 18.04 and Ubuntu 20.04) and Windows (10 and 11). We treat the data of HUAWEI KLV-W19L with Ubuntu 18.04 as the training set and test under the other three hosts. The ASRs and FARs are shown in Fig 15. It can be observed that all the ASRs are higher than 99%. Even when the training set and test set are from different hosts with varying OS types, the ASR is not affected. Simultaneously, the FARs under all testing conditions are very low. These results demonstrate that *DiskPrint*'s authentication performance is independent of both the host model and operating system. Consequently, the same fingerprints can be transferred between different hosts and are not affected by the OSes either.

■ **Effect of Occupation Rate.** In practice, the disk may store different amounts of data. To study if the occupation rate impacts *DiskPrint*'s performance, we vary it from 0% to 80% in step of 20% to collect fingerprint signals. Then, we use the data of 0% as the training set and test under other occupation rates. The ASRs and FARs are shown in Fig. 16. It can be seen that the ASR does not vary obviously with the occupation rate. Although the FAR changes when the occupation rate increases, it is always lower than 1.7%. This indicates that the occupation rate has negligible impacts on *DiskPrint*'s performance. Users only need to register under one occupation rate.

■ **Effect of Application.** It would be common that a number of applications and the stimulation program of *DiskPrint* are co-running on the host. In this experiment, we explore the impact of such concurrent running. Similar to the experimental methodology in [7], we first select five daily-used applications: WeChat, Microsoft Word, Google Chrome, YouTube, and MATLAB. Then, we train the classifiers using EM signals with no other application and test them using EM signals with one of the aforementioned applications, respectively. The ASRs and FARs are shown in Fig. 17. It can be found that the ASRs/FARs maintain very high/low when the application changes. It is apparent that *DiskPrint* is also robust against the application changes.

■ **Cross-position Performance.** It is hard to make the sensor antenna's position in each authentication attempt exactly the same as that in the registration phase. So, we assess the cross-position performance by varying the position from -20 mm to 20 mm in step of 5 mm (marked as from -4 to 4, 4 cm in total), as shown in Fig. 22. The training set is collected under 0mm and we calculate the ASRs of other positions. The resulting ASRs in Fig. 18 manifest




Figure 19: DSRs for mimic attacks.




Figure 20: FRRs caused by our replay-resistant methods.




Figure 21: DSRs for replay attacks.

that the best performance can be achieved when the training data and testing data are collected under the same position (i.e., 0 mm). Once the antenna moves, the ASR will change slightly. But as long as the antenna is within [-20mm, 20mm], the ASR still keeps high. It is easy for a person to let the antenna’s position lie in this region, indicating the good usability of *DiskPrint*.

■ **Cross-time Evaluation.** In this experiment, we treat the fingerprint collected on Day 1 and Day 2 (with two months in between) as training and test sets, respectively. Note that during this period, these disks are in normal use. The resulting ASR/FAR is higher/smaller than 99.5%/1.5%. This proves that the performance of *DiskPrint* does not deteriorate over time.

5.4 Security Study

In this part, we evaluate the security of *DiskPrint* against mimic attack and replay attack.

■ **Mimic Attack Detection.** To strictly assess the capability of *DiskPrint* in defending against mimic attacks, we consider the case that an attacker mimics a target disk (i.e., a legal disk) using a disk that has the same model. We randomly select ten pairs of disks. In each pair, one disk is the victim and the other is the attacker. For each victim, we launch 100 times of attacks and calculate the DSR. The results are shown in Fig. 19. It can be observed that most of the DSRs are 100%. Even the lowest DSR is as high as 98%. This outcome is reasonable as the EM signals emitted from two disks with the same model are also highly distinguishable. Hence, *DiskPrint* is able to effectively defend against mimic attacks.

■ **Replay Attack Detection.** In *DiskPrint*, a replay attack is detected by validating the random bit sequence. If *DiskPrint* wrongly recognizes some bits, it may reject a legal authentication request. Therefore, we first evaluate the FRR of our bit sequence extraction method. Specifically, we randomly select ten disks and perform 100 times of authentication attempts for each disk. The FRRs caused by the wrong bit sequence are shown in Fig. 20. It can be seen that there are no obvious FRR differences among disks. Meanwhile, the

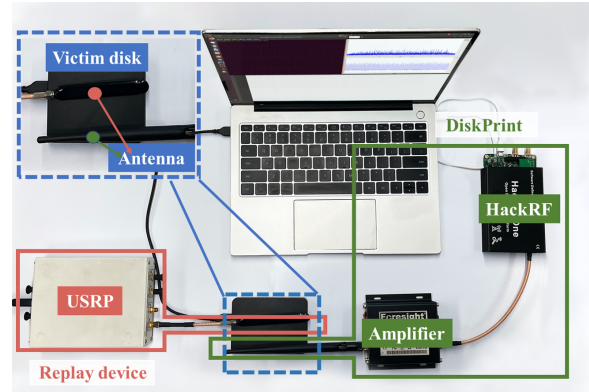


Figure 23: Experiment setup for replay attack.

FRRs of most disks are 0%. This indicates that our replay-resistant method is fair to different disks and has outstanding usability. Then, we evaluate *DiskPrint*’s ability of defending against replay attacks by simulating real-world attacks. As shown in Fig 23, a USRP B210 is utilized to maliciously record the EM signal leaked in the authentication attempt. The antenna is placed on the victim disk surface. The recorded signals are replayed and received by *DiskPrint*. Also, we randomly select 10 disks as the attack targets. The recorded signals of each disk are manipulated according to the protocol of *DiskPrint* in advance. For each target disk, we launch 100 times of replay attacks. The results in Fig. 21 demonstrate that our replay-resistant method successfully detects all replayed signals. Thus, our replay-resistant method can effectively prevent *DiskPrint* from replay attacks.

5.5 Latency

The workflow of *DiskPrint* comprises two phases: registration and authentication. In this part, we separately evaluate the latency of these two phases. *DiskPrint* leverages a supervised learning algorithm to determine the authenticity of disk ID. The time spent collecting the training set accounts for the bulk of the total registration time. As stated in Sec. 5.2, collecting 60 samples costs less than 1.7 minutes. Taking the latency for classifier training into consideration, the total registration time is less than two minutes. The time spent processing an authentication request is mainly composed of five components: signal collection, replay detection, signal enhancement, feature extraction, and disk classification. With a 3.2GHz i7-8700 CPU, the latency of these components is 4.14, 0.24, 0.2, 0.017, and 0.007 seconds, respectively. The low total latency, ~4.6 seconds, indicates the good real-time performance of *DiskPrint*.




Figure 22: Experiment setting for cross-position performance evaluation.

6 DISCUSSION

We also discuss the possibility of two more challenging attacks: truncate attack and man-in-the-middle attack (MITM).

Truncate Attack. In such an attack, the attacker first eavesdrops on a legal authentication attempt (of Disk A) and then employs an illegal disk (Disk B) to initiate an authentication request. In this process, the attacker lets *DiskPrint* normally receive the preamble signal. But when Disk B starts to generate its fingerprint signal, the attacker blocks it with a metal casing and replays the pre-recorded one (of Disk A) to *DiskPrint*'s sensor. In this way, *DiskPrint* could falsely recognize Disk B as A. While theoretically feasible, this attack is significantly difficult to mount in reality. In effect, the signal induced from a disk is continuous (as shown in Fig. 8(b)). But in a truncate attack, the spliced signal is very likely to be noncontinuous. Even if the attack immediately replays the fingerprint signal when s/he detects the end of the preamble signal, the time lag between the detection and reception of replayed signal would cause discontinuity. Furthermore, in this attack, as the preamble signal and fingerprint signal are from different devices and positions, the signal strengths of them are also distinct with high probability. Therefore, it is easy to mitigate the truncate attack, e.g., via discontinuity detection and strength consistency detection.

MITM. Before performing such an attack, the attacker first eavesdrops on a legal authentication process to record the preamble signal and fingerprint signal (of Disk A). Then, the attacker uses an illegal disk (Disk B) to initiate an interaction request. *DiskPrint* will let the stimulation program issue instructions to the disk for bit sequence embedding (i.e., write caching state switching) and fingerprint generation. Once the attacker intercepts these instructions, s/he can “weave” fake preamble signal and fingerprint signal (of Disk A) based on the previous recording, and transmit them to *DiskPrint*'s sensor (while shielding Disk B's emanation). In this case, our replay-resistant mechanism will be deceived and *DiskPrint* may mistake Disk B for A. It can be found that this attack is technically sophisticated and demands special equipment that is able to record, transmit, and craft the fake signal in a timely way. In fact, it is significantly difficult and costly to realize such equipment. On the other hand, this attack is detectable. When performing MITMs, the interception and analysis of the instructions, as well as the editing of the forgery signal are time-consuming. While in normal authentication, *DiskPrint*'s sensor instantaneously receives the authentic signal from the disk. Thus, the time lag between instruction issuing and fake signal reception is significantly greater than that between instruction issuing and authentic signal reception. Again, it would be easy to resist MITM by leveraging the time lag differences, e.g., via setting a 37-millisecond threshold that is larger than the time lags of 99% legal authentication attempts in our experiments.

7 RELATED WORK

Device fingerprinting. Fingerprint-based device authentication can be categorized into software-based and hardware-based solutions. The former ones use browser properties [54], MAC address [15], the feature of the device driver network [14], traffic pattern [33], and etc., as fingerprints. However, they can be easily modified or spoofed [31, 39, 44]. In contrast, hardware-based methods exploit the device's physical components or properties as the

fingerprint, which remain stable throughout the device's lifecycle. Hardware-based fingerprints can be obtained from multiple sources, such as the device's built-in sensors [27, 55], PUF [13, 43], magnetic field leakage [7], and wireless radios [46]. However, sensor-based and PUF methods are limited in universality due to their reliance on specific built-in sensing modules or circuits [13, 27, 55], while magnetic field- and wireless radio-based systems are prone to replay attacks [11, 18]. In this paper, we propose *DiskPrint* which uses EM emanation to wirelessly probe the disk fingerprint and introduce a device-agnostic replay-resistant method for enhanced security.

To better demonstrate *DiskPrint*'s superiority, we compare it with four state-of-the-art USB device fingerprinting systems. As shown in Table 2, the comparison is conducted in terms of non-intrusiveness, cross-host universality, resilience to replay attack, modification on host, and resistance to electromagnetic interference. To be specific, *DeviceVeil* [43] accurately identifies USB devices using PUF. However, it requires invasive hardware modifications of individual USB peripherals, rendering it incompatible with existing legacy devices and severely limiting its practicality. *Time-Print* [10] authenticates USB flash drives by exploiting distinctive timing differences in read operations. Nevertheless, *Time-Print* requires modification of the host driver, and the extracted fingerprints are specific to the host model. *PowerID* [41] fingerprints USB peripherals based on their power consumption under various working conditions. However, the fingerprints extracted by this method may be host model-specific due to many factors (e.g., variations in the power delivery of the host's USB port). These variations can alter the power, potentially limiting fingerprint universality across different host models. *Magneto* [18] leverages the unintentional electromagnetic emissions during host boot operations to fingerprint USB flash drives. However, *Magneto* flaws in the following aspects: 1) the fingerprints extracted from different hosts are distinct, i.e., host-specific; 2) the system is susceptible to replay attacks and lacks a corresponding defense mechanism; 3) the system assumes that the tested peripheral is in an electromagnetic-safe zone, which may be challenging in real-world scenarios. In contrast, *DiskPrint* can work in environments with electromagnetic interference (e.g., laboratories). Moreover, the fingerprints extracted by *DiskPrint* are universal across different host models. Importantly, *DiskPrint* is non-intrusive. It requires no modifications to the disks or hosts. Meanwhile, it demonstrates resilience to various attacks, including replay and mimic attacks.

EM side-channel. The EM side-channel has been extensively studied for both benign and malicious purposes. For benign ones, researchers exploit the leaked EM signals to realize plenty of sensing and communication tasks, including memory profiling [38], hardware/software attestation [18, 18], malware detection [16], covert communication [40], user fingerprinting [12, 53], etc. However, EM side-channel can also be harnessed as an attack vector for privacy mining, e.g., notorious TEMPEST attacks [21, 25] and wireless charging eavesdropping [24], and magnetic field-based computer privacy inference [3]. Despite the much effort devoted to exploiting the EM side-channel, the EM emanation of a disk is overlooked. In this work, we propose *DiskPrint*, to our best knowledge, the first EM side-channel-based disk fingerprinting system that can detect replay attacks by controlling leaked EM signal patterns.

Table 2: Comparison with previous systems. (“———” means that this work does not consider this aspect.)

System	Non-intrusiveness	Cross-hosts	Replay Attack	Without Host Modification	Electromagnetic Interference
		Universality	Resilience		Resistance
<i>DeviceVeil</i> [43]	×	———	———	✓	✓
<i>MAGNETO</i> [18]	✓	×	×	✓	×
<i>Time-Print</i> [10]	✓	×	———	×	✓
<i>PowerID</i> [41]	✓	Unknown	———	✓	✓
<i>DiskPrint</i>	✓	✓	✓	✓	✓

8 CONCLUSION

This paper presents *DiskPrint*, an EM emanation-based replay-resistant disk authentication system. We first demonstrate the feasibility of extracting disk fingerprint from its EM leakage by rigorously modeling the data writing process. Then, we design a series of signal enhancement techniques to improve the signal’s quality. With several carefully selected features, *DiskPrint* achieves accurate disk authentication via an ensemble learning technique. To further secure *DiskPrint*, we design a device-independent anti-replay method by introducing randomness into the EM emanation. Extensive experiments with 60 disks show that *DiskPrint* can achieve 99%+ ASR and 1.5% FAR. Meanwhile, *DiskPrint* is robust against many kinds of context changes. Security study demonstrates that *DiskPrint* is capable of defending against intractable mimic and replay attacks.

ACKNOWLEDGMENTS

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. 2024C03287, and the Postdoctoral Fellowship Program of CPSF under Grant Number GZC20241488.

REFERENCES

- [1] 3dBi SubG Antenna. 2023. 3dBi Antenna. <https://store.rakwireless.com/products/3-dbi-lora-antenna>.
- [2] Zhongjie Ba, Sixu Piao, Xinwen Fu, Dimitrios Koutsonikolas, Aziz Mohaisen, and Kui Ren. 2018. ABC: Enabling Smartphone Authentication with Built-in Camera. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*. The Internet Society. https://www.ndss-symposium.org/wp-content/uploads/2018/02/ndss2018_03B-3_Ba_paper.pdf
- [3] Sebastian Biedermann, Stefan Katzenbeisser, and Jakub Szefer. 2015. Hard drive side-channel attacks using smartphone magnetic field sensors. In *International Conference on Financial Cryptography and Data Security*. Springer, 489–496.
- [4] Encyclopaedia Britannica. 2023. Maxwell’s equation. <https://www.britannica.com/science/Maxwells-equations>.
- [5] Robert Callan, Alenka Zajić, and Milos Prvulovic. 2015. FASE: Finding amplitude-modulated side-channel emanations. *ACM SIGARCH Computer Architecture News* 43, 3S (2015), 592–603.
- [6] CHRISTIAN CAWLEY. 2018. 5 Ways Data Can Be Stolen From Your PC or Network. <https://www.makeuseof.com/tag/how-data-gets-stolen/>.
- [7] Yushi Cheng, Xiaoyu Ji, Juchuan Zhang, Wenyan Xu, and Yi-Chao Chen. 2019. DeMiCPU: Device Fingerprinting with Magnetic Signals Radiated by CPU. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*.
- [8] COM-POWER. 2023. EMI And EMC Testing. <https://www.com-power.com/blog/why-emi-emc-testing-necessary>.
- [9] Casey Crane. 2021. What Is a Device Certificate? Device Certificates Explained. <https://www.thesslstore.com/blog/what-is-a-device-certificate-device-certificates-explained/>.
- [10] Patrick Cronin, Xing Gao, Haining Wang, and Chase Cotton. 2022. Time-Print: Authenticating USB Flash Drives with Novel Timing Fingerprints. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*.
- [11] Boris Danev, Davide Zanetti, and Srdjan Capkun. 2012. On physical-layer identification of wireless devices. *Comput. Surveys* 45, 1 (2012), 6:1–6:29.
- [12] Dian Ding, Lanqing Yang, Yi-Chao Chen, and Guangtao Xue. 2021. Leakage or Identification: Behavior-irrelevant User Identification Leveraging Leakage Current on Laptops. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 4 (2021), 152:1–152:23. <https://doi.org/10.1145/3494984>
- [13] Mohammad Ebrahimabadi, Mohamed Younis, Wassila Lalouani, and Naghme Karimi. 2021. A novel modeling-attack resilient arbiter-PUF design. In *Proceedings of the International Conference on VLSI Design and 2021 20th International Conference on Embedded Systems (VLSID)*.
- [14] Jason Franklin and Damon McCoy. 2006. Passive Data Link Layer 802.11 Wireless Device Driver Fingerprinting. In *Proceedings of the USENIX Security Symposium*.
- [15] Fanglu Guo and Tzi-cker Chiueh. 2005. Sequence Number-Based MAC Address Spoof Detection. In *Proceedings of the International Symposium on Research in Attacks, Intrusions and Defenses (RAID)*.
- [16] Yi Han, Sriharsha Etigowni, Hua Liu, Saman A. Zonouz, and Athina P. Petropulu. 2017. Watch Me, but Don’t Touch Me! Contactless Control Flow Monitoring via Electromagnetic Emanations. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*.
- [17] Computer Hope. 2022. Serial number. <https://www.computerhope.com/jargon/v/vendorid.htm>.
- [18] Omar Adel Ibrahim, Savio Sciancalepore, Gabriele Oligeri, and Roberto Di Pietro. 2021. MAGNETO: Fingerprinting USB Flash Drives via Unintentional Magnetic Emissions. *ACM Transactions on Embedded Computing Systems (TECS)* 20, 1 (2021), 8:1–8:26.
- [19] Foresight Intelligent. 2023. Foresight low noise amplifier (FST-RFAMP02). <https://m.tb.cn/h.UrzGmlQ?tk=ipnmd9AHgeC>.
- [20] Rich Kolko. 2021. New, out-of-the-box, external hard drives could steal your data. <https://winknews.com/2021/01/11/new-out-of-the-box-external-hard-drives-could-steal-your-data/>.
- [21] Ho Seong Lee, Dong Hoon Choi, Kyuhong Sim, and Jong-Gwan Yook. 2018. Information recovery using electromagnetic emanations from display devices under realistic environment. *IEEE Transactions on Electromagnetic Compatibility* 61, 4 (2018), 1098–1106.
- [22] Jianwei Liu, Wenfan Song, Leming Shen, Jinsong Han, Xian Xu, and Kui Ren. 2021. MandiPass: Secure and Usable User Authentication via Earphone IMU. In *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS)*.
- [23] Jianwei Liu, Xiang Zou, Jinsong Han, Feng Lin, and Kui Ren. 2020. BioDraw: Reliable Multi-Factor User Authentication with One Single Finger Swipe. In *Proceedings of the IEEE/ACM International Symposium on Quality of Service (IWQoS)*.
- [24] Jianwei Liu, Xiang Zou, Leqi Zhao, Yusheng Tao, Sideng Hu, Jinsong Han, and Kui Ren. 2022. Privacy Leakage in Wireless Charging. *IEEE Transactions on Dependable and Secure Computing (TDSC)* (2022).
- [25] Zhuoran Liu, Niels Samwel, Leo Weissbart, Zhengyu Zhao, Dirk Laurent, Lejla Batina, and Martha A. Larson. 2021. Screen Gleaning: A Screen Reading TEMPEST Attack on Mobile Devices Exploiting an Electromagnetic Side Channel. In *Proceedings of the Annual Network and Distributed System Security Symposium*

- (NDSS).
- [26] RVE Lovelace, JM Sutton, and EE Salpeter. 1969. Digital search methods for pulsars. *Nature* 222 (1969), 231–233.
 - [27] Dominik Christian Maier, Henrik Erb, Patrick Mullan, and Vincent Haupt. 2020. Camera Fingerprinting Authentication Revisited. In *Proceedings of the International Symposium on Research in Attacks, Intrusions and Defenses (RAID)*.
 - [28] Len Luet Ng, Kim Ho Yeap, Magdalene Wan Ching Goh, and Veerendra Dakulagi. 2022. Power Consumption in CMOS Circuits. In *Electromagnetic Field in Advancing Science and Technology*. IntechOpen.
 - [29] Kim Thuat Nguyen, Maryline Laurent, and Nouha Oualha. 2015. Survey on secure communication protocols for the Internet of Things. *Ad Hoc Networks* 32 (2015), 17–31.
 - [30] Michael Ossmann. 2021. HackRF One. <https://greatscottgadgets.com/hackrf/one/>.
 - [31] Pierluigi Paganini. 2014. How cybercrime exploits digital certificates. <https://resources.infosecinstitute.com/topic/cybercrime-exploits-digital-certificates/>.
 - [32] Hong Pan, Mohsen Azimi, Guoqing Gui, Fei Yan, and Zhibin Lin. 2018. Vibration-based support vector machine for structural health monitoring. In *Experimental Vibration Analysis for Civil Structures: Testing, Sensing, Monitoring, and Control 7*. Springer, 167–178.
 - [33] Jeffrey Pang, Ben Greenstein, Ramakrishna Gummadi, Srinivasan Seshan, and David Wetherall. 2007. 802.11 user fingerprinting. In *Proceedings of the Annual International Conference on Mobile Computing and Networking (MobiCom)*.
 - [34] Jeong Hoon Park, Sang Jin Kim, Jeong Ho Lee, Chang Joon Yoo, Hyo Jin Kang, Byung Cheol Lee, and Jae Goan Jeong. 2016. Effect of CoSi₂ formation process on CMOS transistor electrical properties for sub-100-nm memory applications. *ECS Journal of Solid State Science and Technology* 5, 5 (2016), P264.
 - [35] RANDY. 2022. Can an External Hard Drive Get A Virus? <https://whatsabyte.com/can-external-hard-drive-get-virus/>.
 - [36] RebellionResearch. 2023. What are the advantages and disadvantages of random forest? <https://www.rebellionresearch.com/what-are-the-advantages-and-disadvantages-of-random-forest>.
 - [37] Doug Rollins. 2012. An Overview of SSD Write Caching. https://static.spiceworks.com/attachments/post/0013/5918/ssd_write_caching_tech_brief_lo.pdf.
 - [38] Nader Sehatbakhsh, Alireza Nazari, Alenka Zajic, and Milos Prvulovic. 2016. Spectral profiling: Observer-effect-free profiling by monitoring EM emanations. In *Proceedings of the Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*.
 - [39] Sudip Sengupta. 2022. What is Password Attack. <https://crashtest-security.com/password-attack/>.
 - [40] Cheng Shen, Tian Liu, Jun Huang, and Rui Tan. 2021. When LoRa Meets EMR: Electromagnetic Covert Channels Can Be Super Resilient. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*.
 - [41] Ricardo Spolaor, Hao Liu, Federico Turrin, Mauro Conti, and Xiuzhen Cheng. 2023. Plug and Power: Fingerprinting USB Powered Peripherals via Power Side-channel. In *IEEE INFOCOM 2023 - IEEE Conference on Computer Communications, New York City, NY, USA, May 17-20, 2023*. IEEE, 1–10. <https://doi.org/10.1109/INFOCOM53939.2023.10229048>
 - [42] Hyuk Sun, Kazuki Sobue, Koichi Hamashita, Tejasvi Anand, and Un-Ku Moon. 2019. A 951-fs rms period jitter 3.2% modulation range in-band modulation spread-spectrum clock generator. *IEEE Journal of Solid-State Circuits* 55, 2 (2019), 426–438.
 - [43] Kuniyasu Suzuki, Yohei Hori, Kazukuni Kobara, and Mohammad Mannan. 2019. DeviceVeil: Robust Authentication for Individual USB Devices Using Physical Unclonable Functions. In *49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2019, Portland, OR, USA, June 24-27, 2019*. IEEE, 302–314. <https://doi.org/10.1109/DSN.2019.00041>
 - [44] TheHack.com.br. 2019. Hackers Attack Uber Vendor Leaking Some User Credentials. <https://blog.axur.com/en/credential-leaks-how-they-work-and-why-you-should-be-concerned>.
 - [45] Vijay. 2023. Top 11 Best External Hard Disk. <https://www.softwaretestinghelp.com/best-external-hard-disk/>.
 - [46] Ge Wang, Haofan Cai, Chen Qian, Jinsong Han, Xin Li, Han Ding, and Jizhong Zhao. 2018. Towards Replay-resilient RFID Authentication. In *Proceedings of the Annual International Conference on Mobile Computing and Networking (MobiCom)*.
 - [47] Ge Wang, Haofan Cai, Chen Qian, Jinsong Han, Shouqian Shi, Xin Li, Han Ding, Wei Xi, and Jizhong Zhao. 2020. Hu-Fu: Replay-Resilient RFID Authentication. *IEEE/ACM Transactions on Networking (TON)* 28, 2 (2020), 547–560.
 - [48] Wikipedia. 2023. Median Absolute Deviation. https://en.wikipedia.org/wiki/Median_absolute_deviation.
 - [49] Wikipedia. 2023. Microprocessor chronology. https://en.wikipedia.org/wiki/Microprocessor_chronology.
 - [50] Wikipedia. 2023. Password. <https://en.wikipedia.org/wiki/Password>.
 - [51] Wikipedia. 2023. Serial number. https://en.wikipedia.org/wiki/Serial_number.
 - [52] Weiye Xu, Wenfan Song, Jianwei Liu, Yajie Liu, Xin Cui, Yuanqing Zheng, Jinsong Han, Xinhuai Wang, and Kui Ren. 2022. Mask does not matter: anti-spoofing face authentication using mmWave without on-site registration. In *Proceedings of the Annual International Conference on Mobile Computing and Networking (Mobicom)*.
 - [53] Lanqing Yang, Yi-Chao Chen, Hao Pan, Dian Ding, Guangtao Xue, Linghe Kong, Jiadi Yu, and Minglu Li. 2020. MagPrint: Deep Learning Based User Fingerprinting Using Electromagnetic Signals. In *39th IEEE Conference on Computer Communications, INFOCOM 2020, Toronto, ON, Canada, July 6-9, 2020*. IEEE, 696–705. <https://doi.org/10.1109/INFOCOM41043.2020.9155534>
 - [54] Ting-Fang Yen, Yinglian Xie, Fang Yu, Roger Peng Yu, and Martín Abadi. 2012. Host Fingerprinting and Tracking on the Web: Privacy and Security Implications. In *Proceedings of the Annual Network and Distributed System Security Symposium (NDSS)*.
 - [55] Zhe Zhou, Wenrui Diao, Xiangyu Liu, and Kehuan Zhang. 2014. Acoustic Fingerprinting Revisited: Generate Stable Device ID Stealthily with Inaudible Sound. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*.