Turning GPU into an FM Radio: A Practical Data Exfiltration Framework from Air-gapped Systems

Rui Xiao[1,2], Sibo Feng[3], and Jinsong Han[3*]
[1]Shanghai University of Finance and Economics
[2]MoE Key Laboratory of Interdisciplinary Research of Computation and Economics, [3]Zhejiang University
xiaorui@sufe.edu.cn, sibofeng@zju.edu.cn, hanjinsong@zju.edu.cn

***Abstract*—In the field of cybersecurity, it has been widely believed that air-gapped devices, disconnected from the Internet, offer strong protection against data breaches. However, recent research has challenged this assumption by demonstrating the feasibility of exfiltrating data from air-gapped devices using *covert channels*, such as modulating thermal, acoustic, or electromagnetic (EM) emanations. Existing approaches, however, face key limitations, including low bit rates, line-of-sight requirements, and susceptibility to EM shielding, limiting their practical deployments. To this end, we propose *MagWhisper*, a novel and practical data exfiltration framework that exploits GPU magnetic leakages for a covert channel. The non-radiative magnetic signals can penetrate various materials, including metal, enabling *MagWhisper*'s non-line-of-sight and shielded operation. We design a high-frequency frequency-shift keying modulation scheme over GPU magnetic signals, thus establishing a high-rate and robust channel. We further adopt a GPU-agnostic design to enhance *MagWhisper*'s generalizability. A proof-of-concept system is implemented to demonstrate *MagWhisper*'s feasibility. Our evaluation on three heterogeneous GPUs shows that *MagWhisper* achieves high bit rates exceeding 133 bps with a bit error rate (BER) lower than 1.2%. Furthermore, real-world case studies confirm that *MagWhisper* remains effective in non-line-of-sight scenarios and can even bypass EM shielding within a Faraday cage.**

I. Introduction

In today's digital landscape, data security has emerged as a top priority. To combat data theft, a range of network security measures is implemented, such as traffic monitors and firewalls. Among them, *air gaps* stands as a foundational approach, wherein the computing devices to protect are purposefully isolated from external networks, thereby safeguarding sensitive data within the confines of the internal network [1]. This approach has been widely adopted in critical environments such as academic institutions, financial organizations, and government entities to protect their most valuable data. More recently, companies such as OpenAI and Microsoft have also reported using air-gapped setups to secure their critical AI infrastructures [2].

However, while air-gapped infrastructures were regarded as secure, researchers and attackers have demonstrated the feasibility of constructing *covert channels* to breach air gaps and realize unauthorized data exfiltration [3], [4], [5]. These methods do not rely on conventional communication protocols (e.g., WiFi) and instead exploit *nontraditional artifacts* for communication. For instance, researchers have demonstrated covert data transmission utilizing cooling fans, where variations in fan noise represent binary information (e.g., high noise for "1" and low noise for "0") [4]. The data can then be received using a microphone. Similar techniques involve leveraging thermal emanation, acoustics, or electromagnetic (EM) emanation [3], [5], [6]. However, their limited performance and strict deployment constraints hinder practical adoption. Built on unintentional physical artifacts, these channels typically offer very low throughput, e.g., 10 bits per second (bps) [3], [4], [5], [7], [8]. Moreover, their requirement of line-of-sight (LoS) transmission and susceptibility to electromagnetic (EM) shielding further render them impractical in real-world scenarios.

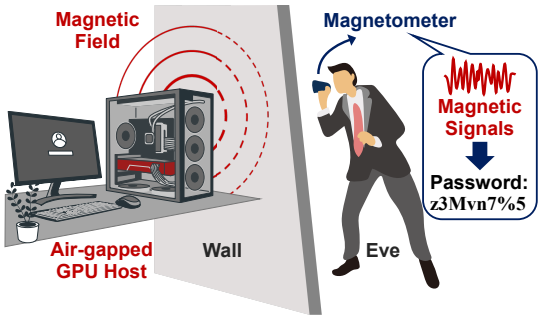

Fig. 1: Figure depicts the attack scenario of *MagWhisper*. The victim GPU host has no network interfaces, i.e., air-gapped. With *MagWhisper*, the attacker (Eve) can establish a magnetic covert channel to steal confidential data even behind a wall.

In light of this, we pose the following question: *Is it possible to design a more sophisticated technique that allows an attacker to exfiltrate data at a higher data rate and remains effective under non-line-of-sight (NLoS) and EM shielding scenarios?* To this end, we introduce *MagWhisper*, a novel data exfiltration framework that capitalizes on graphic processing units (GPUs), which are increasingly prevalent due to the rise of AI infrastructures. The core idea behind *MagWhisper* lies in the fact that GPUs, being highly power-intensive, exhibit *potent physical signal leakages*. Among them, we are particularly interested in the *non-radiative magnetic signals*, which can traverse walls and remain resilient even when

*Jinsong Han is the corresponding author.

EM shielding is applied, making them an ideal medium for covert communication. As depicted in Fig. 1, an attacker can discreetly receive credential data from an air-gapped host using a concealed magnetometer, e.g., positioned behind a wall.

These promising characteristics of GPU magnetic emanation might suggest that building a covert channel would be straightforward – by employing amplitude-shift-keying (ASK) modulation that toggles GPU states (e.g., active for "1" and idle for "0"). However, this design faces two primary challenges. The first challenge is the high level of noise caused by external magnetic interference from nearby electronic components like CPUs and power supplies. These sources introduce *amplitude shifts* and pulse-like noises into the signals. Furthermore, the magnetic signal attenuates with distance, making it increasingly difficult to separate the encoded amplitude from background interference. Consequently, the ASK-based solution is unfeasible. An alternative is frequency-shift keying (FSK) modulation, which is more resilient to amplitude-shift interference [9]. However, existing FSK designs rely on toggling between idle and active GPU functions – a form of "function-level" switching that typically operates below 13 Hz, thereby severely limiting the data rate [7], [10]. The second challenge stems from GPU heterogeneity. Different GPU models exhibit distinct hardware specifications, resulting in varying magnetic leakage signals. This diversity complicates the development of a robust demodulator that can generalize across devices.

To address the first challenge of external interference while maintaining a high data rate, we design an innovative high-frequency FSK implementation on the transmitter side. Instead of relying on coarse-grained "function-level" switching, *MagWhisper* employs fine-grained "operation-level" switching between arithmetic and memory operations within a function. This approach yields stable, high-frequency magnetic emissions in the kilohertz range, significantly boosting the data rate. On the receiver side, we apply meticulous signal processing techniques, including pulse-noise removal and signal detrending, to further enhance the transmission robustness. To overcome the second challenge of GPU dependency, we draw inspiration from channel estimation techniques commonly employed in RF-based systems. By embedding a known preamble and extracting GPU-specific parameters during reception, *MagWhisper* achieves a GPU-agnostic approach, enhancing the adaptability and robustness of the covert channel across heterogeneous hardware.

MagWhisper eliminates the need for bulky or costly receiver hardware required by other covert channels, such as mmWave radars [5]. In our proof-of-concept implementation, we use a low-cost magnetometer (around $3) for signal acquisition and a Raspberry Pi for post-processing, resulting in a compact overall form factor. Since *MagWhisper* operates effectively in NLoS conditions and can penetrate shielding, the receiver can be discreetly concealed, for instance, inside a drawer or behind a wall (§VII-D), significantly enhancing its stealthiness.

We comprehensively evaluate *MagWhisper*'s performance on three heterogeneous GPUs under various conditions, including different wall obstructions (e.g., metal), sensor locations, and GPU workloads. Additionally, we conduct two real-world case studies: one involving text transmission through a 25 cm concrete wall and another through a Faraday cage. Overall, *MagWhisper* achieves bit rates exceeding 133 bps with a bit error rate (BER) below 1.2% and demonstrates robustness across a wide range of real-world scenarios.

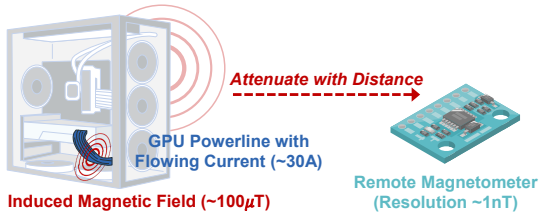

Fig. 2: Figure depicts that the large electric current in the GPU generates a non-radiative magnetic field, which can be received using a remote magnetometer.

Our main contributions are summarized as follows:

- We propose *MagWhisper*, the first covert channel exploiting GPUs' physical artifacts that achieves high data rates and works under NLoS and EM shielding.
- We achieve robust, high-rate FSK modulation on GPUs by leveraging operation-level switching between arithmetic and memory operations. The overall design further enables GPU-agnostic and reliable reception.
- We evaluate *MagWhisper* through comprehensive real-world experiments, demonstrating its high robustness.

II. Background

A. GPU Computation Architecture

Modern computing systems, especially AI infrastructures, widely adopt GPUs as high-performance co-processors for parallel workloads such as deep learning and graphics rendering. GPU computation primarily consists of two types of operations: **arithmetic operations** and **memory operations**. Arithmetic operations perform parallel mathematical computations (e.g., additions, multiplications), leveraging the GPU's ability to execute thousands of threads concurrently – a key feature that enables massive parallelism. In contrast, memory operations manage data access and movement through the dedicated *GPU memory*, which is physically separate from system memory and optimized for high-throughput data exchange. In §II-D, we characterize the distinct magnetic signals produced by these two types of operations, which form the foundation of *MagWhisper*'s high-frequency FSK modulation design.

B. GPU Magnetic Emanation

MagWhisper leverages GPU magnetic emanations as a covert channel. GPUs emit two types of electromagnetic leakages: high-frequency electromagnetic radiation and non-radiative magnetic signals. The former, typically in the MHz range or higher, can be effectively suppressed through EM shielding. In this paper, we mainly focus on the non-radiative magnetic leakages, referred to as *magnetic signals* for simplicity. We also briefly explore the potential of high-frequency EM radiation for long-range covert communication in §VIII.

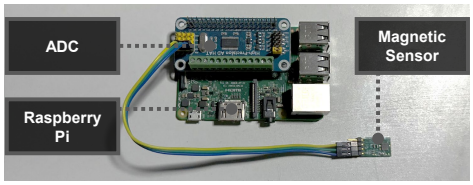

Fig. 3: *MagWhisper*'s magnetic acquisition setup consists of a low-cost magnetic sensor (3 USD), an ADC, and a Raspberry Pi.

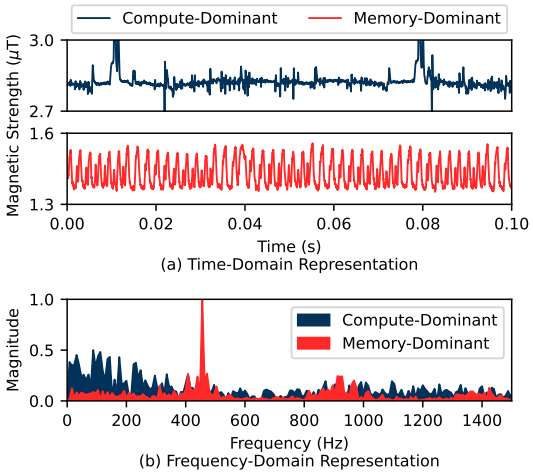

Fig. 4: Figure depicts the magnetic behavior of compute-dominant programs and memory-dominant programs represented in (a) the time domain and (b) the frequency domain.

Fig. 2 shows the cause of GPU magnetic emanation and our sensing methodology. The magnetic signal is generated by the electric current flowing through the power wires connecting the GPU to the power supply unit. According to the Biot-Savart Law [11], the magnetic field generated by the current I can be described as – $\mathbb{B} = \frac{\mu_0 I}{2\pi r}$, where μ_0 is the magnetic constant, and r is the distance from the wire. Consumer-grade GPUs typically have a power consumption ranging from 100W to 450W. With a standard input voltage of 12 volts, this translates to a current of around 8A to 38A, which is sufficient to generate magnetic fields reaching tens of microtesla, well above the sensitivity threshold of the low-cost magnetometer introduced in the next section.

C. Sensing GPU Magnetic Emanation

We adopt fluxgate magnetometers for sensing GPU magnetic emanation due to their superior sensitivity, accuracy, and wider operating range. Specifically, we use the DRV425 fluxgate sensor [12], which offers nanotesla-level resolution at a low cost (approximately $3). Its compact form factor makes it well-suited for covert deployment, allowing the receiver to be easily concealed. The sensor outputs a voltage proportional to magnetic field strength (sensitivity = 5 mV/μT). Fig. 3 illustrates our setup, which includes the DRV425 sensor, an ADS1263 analog-to-digital converter (ADC) for digitizing the magnetic signals at a sampling rate of 19.2 KHz, and a Raspberry Pi for post-processing [13], [14].

D. Characterizing GPU Magnetic Behavior

We now introduce a detailed characterization of GPU magnetic behavior, which is essential for designing an efficient modulation scheme. As discussed in §II-A, GPU programs can be categorized based on the dominant operation type:

- Compute-dominant programs: These programs primarily consist of arithmetic operations, with significantly fewer memory accesses.
- Memory-dominant programs: These programs primarily consist of memory access operations, with relatively less arithmetic computation.

We collect magnetic signals from both program types and present the time-domain traces in Fig. 4(a), along with their frequency-domain representations in Fig. 4(b).

Key Insight. The analysis reveals that compute-dominant programs generate stronger magnetic signals than memory-dominant ones. However, memory-dominant programs exhibit **stable, periodic fluctuations** in the magnetic field, resulting in a distinct and consistent *magnetic frequency*. This critical observation guides the design of *MagWhisper*'s robust high-frequency FSK modulation scheme, which centers on modulating the magnetic frequency of memory-dominant programs.

III. Threat Model

We now outline the attacker's goals and assumptions.

Attacker's goal. The attacker's goal is to execute a stealthy data exfiltration attack by establishing a covert channel leveraging GPU physical leakages from the victim host. This covert channel operates even when there are obstacles between the attacker and the victim host, such as walls, and when the host is EM shielded by Faraday cages or similar measures.

Assumptions. *MagWhisper* builds upon the standard assumptions *commonly used* in conventional covert channels, i.e., requiring the presence of transmitter malware on the victim host [3], [4], [5], [8]. For *MagWhisper*, the planted malware *only requires user-space access* and does not need root privileges. The infection of a computing device can be achieved through various attack vectors, such as social engineering, supply chain attacks, or bundling with legitimate software [15], [16], [17]. Additionally, a malicious actor with access to a shared, corporate, or public computer, like a malicious insider, can install the malware [18]. While *MagWhisper* shares common assumptions with previous covert channel designs, it expands the attacker's goal beyond what many previous works can achieve by leveraging the unique characteristics of GPU magnetic emanation.

IV. Overview of *MagWhisper* Design

We present a concise overview of *MagWhisper*, illustrated in Fig. 5, which consists of a GPU acting as the transmitter and a magnetometer serving as the receiver.

Transmitter: Initially, the credential data, such as passwords, undergoes packetization. The packets are then modulated using FSK modulation and transmitted through the GPU's magnetic driver during the execution of GPU operations. For increased stealthiness, the GPU load can be tuned to reduce its usage.

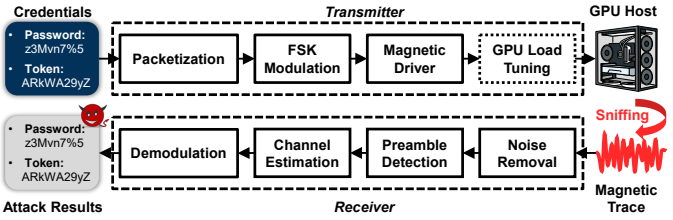

Fig. 5: Figure depicts the overview of *MagWhisper*. The secrets are transmitted through the GPU magnetic channel using FSK modulation and will be demodulated by the receiver with a magnetometer.

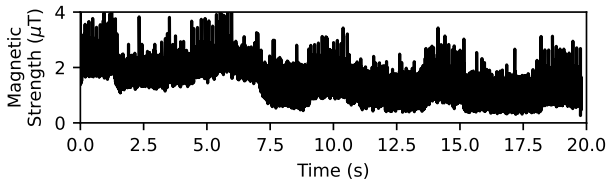

Fig. 6: Figure depicts Challenge I: magnetic channel suffers ambient noises composed of amplitude shifts and pulse-like noises.

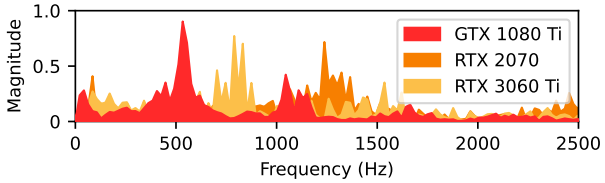

Fig. 7: Figure depicts Challenge II: the magnetic frequency is highly dependent on the GPU model.

Receiver: On the receiver side, the received magnetic signals undergo noise removal to enhance signal quality. Preambles are detected to facilitate channel estimation and symbol segmentation. Finally, the symbols are demodulated, allowing the attacker to extract the transmitted credentials.

Designing *MagWhisper* faces the following challenges –

Challenge I: Interference from Ambient Magnetic Noise. Previous covert channel studies commonly adopt ASK modulation to simplify the design [3], [5], [6]. However, the GPU magnetic channel is particularly susceptible to ambient magnetic noise, especially amplitude shifts. Fig. 6 illustrates the measured magnetic trace at a distance of 30 cm from the host while the GPU is idle, revealing significant fluctuations and jitters with levels as high as 4 μT. These substantial interferences render amplitude-based modulation unfeasible for transmission over the magnetic channel.

An alternative is FSK modulation, which encodes symbols using different frequencies (e.g., 1 kHz for symbol 0 and 1.5 kHz for symbol 1), and is inherently more robust to amplitude fluctuations [9]. However, simply implementing FSK by switching between idle and active functions can only produce low-frequency signals due to the limited switching rate at the function level. This significantly compromises the data rate, limiting it to within 10 bps [7], [10].

Solution: We devise a high-frequency FSK modulation scheme (§V-A) that exploits fine-grained switching between

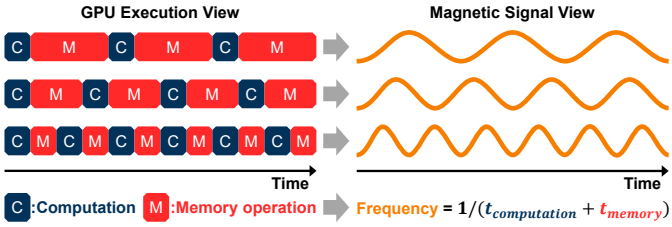

Fig. 8: Figure depicts how the magnetic frequency can be controlled by altering the number of GPU memory and computation operations.

memory and arithmetic operations. This approach enables the generation of kilohertz-range frequencies, significantly improving the data rate. Besides, we apply various noise removal techniques on the receiver side to further enhance the robustness of the demodulation process (§VI-A).

Challenge II: GPU-dependent Magnetic Emanation. When executing the same set of codes, different GPU models exhibit varying magnetic frequencies. As depicted in Fig. 7, the magnetic signals from GTX 1080 Ti and RTX 2070 show frequencies of 550 Hz and 1300 Hz, respectively. This variation arises from architectural heterogeneity across GPU models and presents a significant challenge to achieving robust and consistent demodulation.

Solution: We draw inspiration from channel estimation techniques utilized in RF-based communication systems. By sending a known preamble at the beginning of each packet (§V-B), *MagWhisper* can adaptively demodulated the data (§VI-B), thereby enabling a GPU-agnostic design.

V. *MagWhisper* Transmitter

MagWhisper transmitter aims to utilize the frequency modulation effect of memory-dominant GPU programs to achieve high-frequency FSK modulation on magnetic signals.

A. Generating Different High-frequency Magnetic Signals

The key requirement for FSK modulation is the ability to generate distinct frequencies to represent different symbols. In the case of *MagWhisper*, we control the GPU magnetic frequency by tuning the number of memory reads and computation operations. By alternating between memory accesses and computation operations, we establish a relationship for the magnetic frequency F_{mag}:

$$F_{\text{mag}} = \frac{1}{t_{\text{computation}} + t_{\text{memory}}}, \tag{1}$$

where $t_{\text{computation}}$ and t_{memory} represent the time taken for computation and memory operations, respectively, in an alternating manner. As shown in Fig. 8, by decreasing the time of memory operations t_{memory}, we can increase the magnetic frequency. However, it is crucial to ensure that $t_{\text{memory}} \gg t_{\text{computation}}$ to maintain a memory-dominant program.

To further interpret Equation 1, let us consider a specific GPU G. The equation can be rewritten as:

$$\frac{n_{\text{read}}}{v_{\text{read}}(G)} + \frac{n_{\text{compute}}}{v_{\text{compute}}(G)} = \frac{1}{F_{\text{mag}}}, \tag{2}$$

n_{read}	F_{mag}
10	990 Hz
20	848 Hz
30	720 Hz
40	638 Hz
50	578 Hz
60	510 Hz
70	458 Hz
80	428 Hz

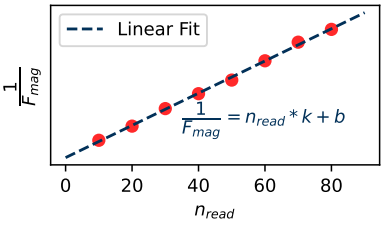

Fig. 9: Figure depicts how the magnetic frequency F_{mag} changes with the number of memory reads n_{read}.

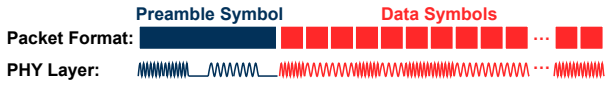

Fig. 10: Figure depicts *MagWhisper*'s packet structure and corresponding physical layer signal.

where n_{read} and n_{compute} denote the numbers of memory reads and computation operations, respectively, and $\upsilon_{\text{read}}(G)$ and $\upsilon_{\text{compute}}(G)$ represent the corresponding execution speeds specific to GPU G.

Fig. 9 presents an experimental validation of Equation 2 by varying n_{read} and measuring the resulting magnetic frequency F_{mag}. Here, each memory read operation corresponds to reading 128 bytes from GPU memory. The results confirm the predicted linear relationship between $\frac{1}{F_{\text{mag}}}$ and n_{read} in Equation 2, demonstrating an effective method for modulating the magnetic frequency during data transmission.

B. Packetization

As demonstrated by Equation 2, different magnetic frequencies F_{mag} can be generated by adjusting n_{read}. However, due to the GPU-dependent nature of υ_{read} and $\upsilon_{compute}$, the resulting frequency on the victim device cannot be precisely controlled. To address this limitation, we introduce a known preamble at the beginning of each packet, allowing the receiver to estimate the transmission's frequency components accurately.

As illustrated in Fig. 10, each packet consists of a preamble symbol followed by consecutive data symbols. The preamble symbol is composed of five "0" symbols and five "1" symbols, separated by a 50ms idle period. The preamble symbols serve a dual purpose. Firstly, they provide a reference for the receiver to accurately estimate the frequency of the magnetic emanation, facilitating channel estimation and synchronization. Secondly, this design enables reliable packet segmentation, allowing the receiver to demodulate the subsequent data symbols. This approach ensures robust and GPU-agnostic communication over the magnetic covert channel.

C. Tunable GPU Load

The *MagWhisper* system incorporates a tunable GPU load design to enhance its stealthiness. This mechanism allows the attacker to control the level of GPU utilization by inserting a configurable dummy period between symbols, during which the GPU remains idle and performs no operations, as illustrated in Fig. 11(a). By adjusting the length of the dummy

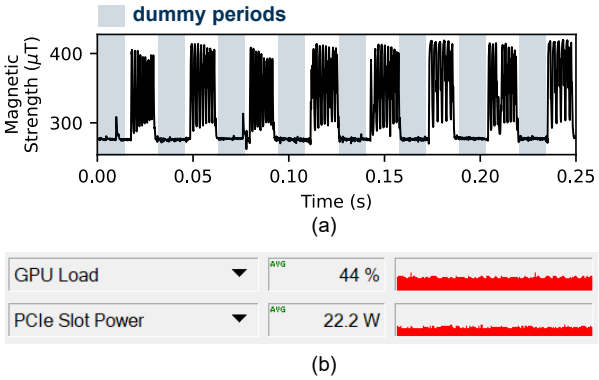

Fig. 11: Figure depicts (a) *MagWhisper*'s tunable GPU load design with dummy periods between symbols, and (b) corresponding GPU load with dummy period insertion.

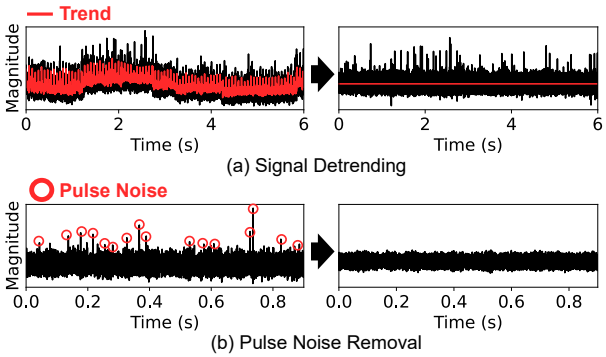

Fig. 12: Figure depicts the effects of (a) Signal Detrending and (b) Pulse Noise Removal steps in the Noise Removal module.

period, the overall GPU load can be fine-tuned within a wide range, from as low as 5% to 100%. As an example, we inserted a 5ms dummy period between symbols and monitored the resulting GPU load and power consumption using GPU-Z [19], a widely used GPU monitoring tool. As shown in Fig. 11(b), the GPU load decreased to 44%, with the PCIe slot power averaging only 22.2W over a one-minute interval. While introducing dummy periods may reduce the effective transmission speed, it enables attackers to trade off performance for stealth. By tuning this parameter, the attacker gains fine-grained control over the system's behavioral footprint, allowing *MagWhisper* to flexibly adapt to various covert communication requirements and deployment constraints.

VI. *MagWhisper* Receiver

We now present *MagWhisper*'s detailed receiver design.

A. Noise Removal

Signal Detrending. As discussed in Challenge I, the magnetic channel is subject to transient amplitude shifts caused by external magnetic noise. To address this issue, as depicted in Fig. 12(a), we remove these amplitude shifts by *detrending* the signal through the subtraction of the moving mean (i.e., the red line) computed over a window of 50 samples (2.6 ms). This method is particularly suited to *MagWhisper*'s FSK

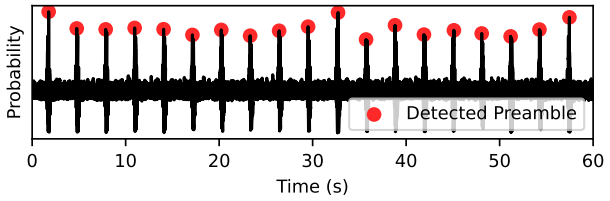

Fig. 13: Figure depicts that *MagWhisper* uses cross-correlation to identify the positions of preambles and segment the packets.

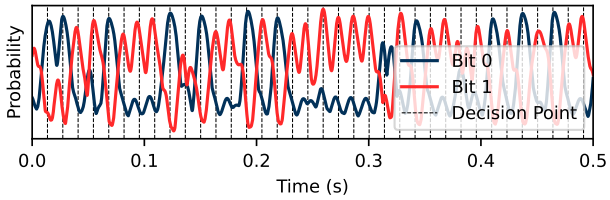

Fig. 14: Figure depicts that *MagWhisper* utilizes the normalized frequency strength as the demodulation metric.

modulation, as it preserves the frequency content essential for demodulation. In contrast, amplitude-based modulation schemes would be adversely impacted by this process, as the amplitude variations carry the encoded information.

Pulse Noise Removal. The magnetic signals may also contain impulsive noise originating from nearby electronic components (e.g., CPUs). To address this, we apply the Hampel filter to detect and suppress outliers in the time series. The filter computes the median m and standard deviation σ within a sliding window of 200 samples. Samples deviating by more than 3σ from the median $m$are classified as outliers and replaced with the median, as illustrated in Fig. 12(b).

B. MagWhisper Demodulation

Preamble Detection. As discussed in §V-B, *MagWhisper* incorporates a fixed preamble at the beginning of each data frame. The preamble includes a 50-ms idle period, making it easily distinguishable from data symbols that contain steady high-frequency components. By identifying the preamble, the transmission frequencies f_0 and f_1, corresponding to Bit-0 and Bit-1, can be determined. Once a single preamble is detected, its fixed structure allows the receiver to accurately identify the start time of each packet using cross-correlation. As illustrated in Fig. 13, packet segmentation is achieved by sliding a window of preamble length across the signal and locating the point with the highest correlation score. This approach ensures precise synchronization and reliable data retrieval by *MagWhisper*.

FFT-based Demodulation. For FSK demodulation, *MagWhisper* utilizes Fast Fourier Transform (FFT)-based frequency analysis [20]. Specifically, it computes the spectral power traces $p_{f_1}(t)$ and $p_{f_2}(t)$ of the transmission frequencies, f_0 and f_1, using a fixed-size sliding window, similar to the short-time Fourier transform. The demodulator then performs *z-normalization* on these two spectral traces [21]. Taking $p_{f_1}(t)$ as an example, the normalized spectral trace $\tilde{p}_{f_1}(t) = \frac{p_{f_1}(t)-\mu}{\sigma}$, where μ and σ are the mean and standard deviation of

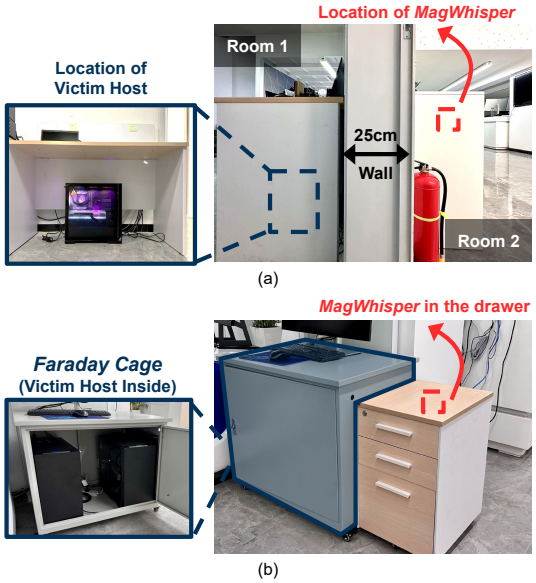

Fig. 15: Figure depicts the two case study setups: (a) cross-room scenario and (b) EM shielding scenario with a Faraday cage.

$p_{f_1}(t)$, respectively. *MagWhisper* then demodulates the data symbols by comparing the normalized power traces at each decision point, as depicted in Fig. 14, allowing *MagWhisper* to accurately retrieve the transmitted data symbols.

VII. Evaluation

We present the evaluation of *MagWhisper* through comprehensive real-world experiments, demonstrating its feasibility.

A. Experiment Setup.

Platform. Our magnetic acquisition setup (Fig. 3) comprises the following components: a DRV425 magnetic-field sensor, an ADS1263 ADC operating at a 19.2 kHz sampling rate, and a Raspberry Pi 4B [12], [13], [14]. The tests are conducted on a computer with a GPU, featuring an AMD Ryzen 5 2600X CPU and running on Windows 10 OS. The transmitter program is implemented using CUDA Toolkit [22].

Data Collection. We assess *MagWhisper*'s performance using three heterogeneous GPU models (GTX 1080 Ti, RTX 2070, and RTX 3060 Ti) released within the past seven years (§VII-B). The default sensing distance is set to 10 cm from the host, and the bit rate is 72 bits per second (bps). We assess *MagWhisper*'s performance under various transmission rates and through-wall scenarios in §VII-B. In §VII-C, we evaluate *MagWhisper*'s robustness concerning different sensing distances and locations, as well as its performance under different sampling rates and frequency spacing. Furthermore, in §VII-D, we investigate two real-world scenarios that are challenging for existing systems, i.e., the cross-room scenario and the EM-shielding scenario (shown in Fig. 15). Additionally, we evaluate the effects of Forward Error Correction (FEC) on *MagWhisper*'s performance. To ensure accurate measurements, for each data point, we average the results from 20 packets, with each packet containing a preamble symbol

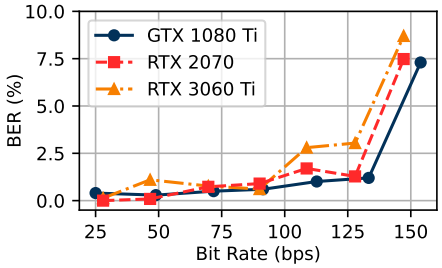

Fig. 16: Figure depicts BERs under different bit rates.

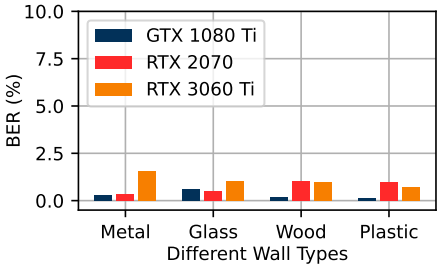

Fig. 17: Figure depicts BERs with different wall types.

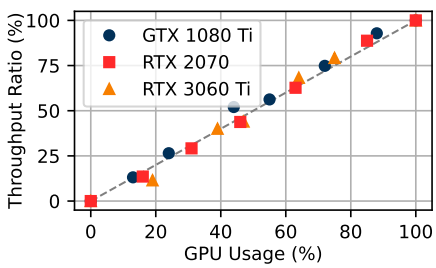

Fig. 18: Figure depicts the effect of the GPU Load Tuning Module.

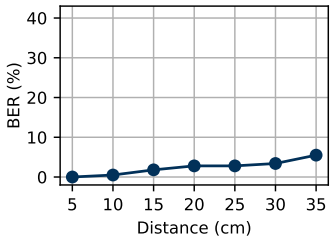

Fig. 19: Figure depicts BERs under different sensing distances.

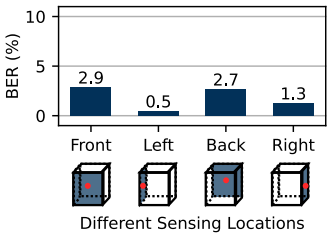

Fig. 20: Figure depicts BERs varying sensing locations.

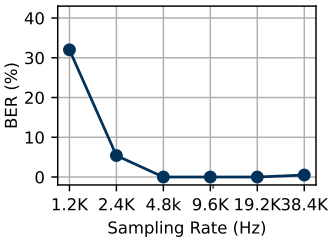

Fig. 21: Figure depicts BERs under different sampling rates.

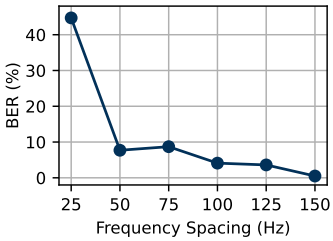

Fig. 22: Figure depicts BERs with different frequency spacing.

and 500 data symbols, resulting in a total of 10,000 symbols for each data point.

Performance Metrics. To assess *MagWhisper*'s effectiveness in data exfiltration attacks, we utilize the metric of **Bit Error Rate (BER)**. BER is computed using the formula $BER = N_{error}/N_{total}$, where N_{error} is the number of bit errors, and N_{total} is the total number of transmitted bits. A lower BER indicates better performance for *MagWhisper*.

B. Overall Communication Performance

We evaluate *MagWhisper*'s communication performance under different transmission rates and through-wall scenarios.

Transmission Capacity. The transmission speed is influenced by the symbol length. A shorter symbol duration results in a higher data rate but also increases vulnerability to noise interference. Fig. 16 illustrates the relationship between bit rate and BER, using six different sets of symbol lengths (corresponding to six different bit rates) on each GPU. Notably, *MagWhisper* achieves lower BER on various GPUs, including 1.2% on the GTX 1080 Ti, 3.0% on the RTX 2070, and 1.3% on the RTX 3060 Ti, all at bit rates higher than 130 bps. When the bit rate is higher than 147 bps, the average BER increases to 7.9% due to due to reduced symbol length. Overall, these results demonstrate the effectiveness of *MagWhisper*'s high-frequency FSK design in enabling reliable, high-throughput transmission. We further compare *MagWhisper* against prior work in §VII-E and highlight its superior performance.

Through-Wall Transmission Performance. One of the key features of *MagWhisper* is its ability to operate as a through-wall covert channel attack. To evaluate this capability, we place various representative materials between the transmitting host and the receiver, i.e., plastics, wood, tempered glass, and metal. As depicted in Fig. 17, *MagWhisper* demonstrates robust performance across all cases, achieving BERs lower than 1.6% with all four occlusion scenarios. This NLoS transmission capability is further validated in our case studies (§VII-D). This enhances the practicality of *MagWhisper* as a covert attack, as the receiver can be concealed behind seemingly innocuous objects or materials while still maintaining successful communication in obstructed environments.

Effect of GPU Load Tuning. To enhance the stealthiness of *MagWhisper*, we introduced a GPU Load Tuning module (§V-C). We now investigate the impact of GPU load tuning on the transmission rate. We define the throughput ratio as the ratio between the bit rates before and after the GPU tuning. We conduct tests with seven different transmission ratios on each GPU, and the results are presented in Fig. 18. The results demonstrate that the GPU usage is approximately linearly proportional to the throughput ratio on all three GPUs. Moreover, this tuning of GPU load is orthogonal to other factors such as bit rate and wall types, making it an easy-to-use technique that improves the stealthiness of *MagWhisper*.

C. Differing Experimental Conditions

We evaluate *MagWhisper*'s performance across several factors using a representative GPU, the NVIDIA GTX 1080 Ti. The reported data is based on an average of 20 packets, each consisting of 500 data symbols.

Impact of Communication Distance. We explore the impact of the distance between the GPU host and the magnetometer on BER. As the distance increases, the strength of the magnetic field decreases following an inverse relationship. The evaluation, shown in Fig. 19, reveals that a BER of 5.5% can still be achieved even at a distance of 35 cm. This demonstrates the robustness of the FSK modulation scheme used in *MagWhisper*, enabling reliable transmission under low

TABLE I: Attack results in the two case study scenarios.

	w/o FEC		with FEC	
	Scenario (a)	Scenario (b)	Scenario (a)	Scenario (b)
BER	2.7%	0.9%	<0.1%	<0.1%

TABLE II: Decoded data without and with FEC.

	Personal Information	**Private Key**
Ground Truth	Date of Birth: January 1, 1900 Address: 123 Main Street, Anytown	MIIFRzCCBC+gAwIBAg IQC2FFwm4d7zP4+7Z1 lACNrzANBgkqhkiG9w 0BAQsFADBH
Received (w/o FEC)	Date of Birth: JaNucr} 1, 1900 Address* 1"3 Main Stseet, Anytown	MIIFRzCCBC+gAwIBAG ISC2FFwm4d7zP4+7Z1 lASNrjANBgkqhkiF9w 0BAQsFADBH
Received (with FEC)	Date of Birth: January 1, 1900 Address: 123 Main Street, Anytown	MIIFRzCCBC+gAwIBAg IQC2FFwm4d7zP4+7Z1 lACNrzANBgkqhkiG9w 0BAQsFADBH

signal-to-noise conditions where traditional amplitude-based modulation schemes typically fail. Moreover, we will show that *MagWhisper* can enable transmission over ***14 meters*** in §VIII with the same *MagWhisper* transmitter.

Impact of Sensing Location. We evaluate *MagWhisper* by placing the sensor at four different locations outside the host device, as depicted in Fig. 20. The results indicate that receiving signals from the left panel of the host device yields the lowest BER compared to other locations. This can be attributed to the proximity of the left panel to the GPU power line, which emits strong magnetic signals. However, it is important to note that the BERs across all locations remain below 2.9%. This evaluation highlights the minimal impact of sensor placement and underscores the usability of *MagWhisper*.

Impact of Sampling Rate. Reducing the required sampling rate of the magnetometer provides several benefits, including lowering hardware complexity and enabling more compact attack devices. In our default configuration, we employ an ADC with a sampling rate of 19.2 KHz. We now vary the sampling rate from 1.2 KHz to 38.4 KHz, as shown in Fig. 21. The results show that *MagWhisper* maintains a BER below 1% when the sampling rate exceeds 4.8 kHz, a rate achievable by low-end ADCs. This demonstrates *MagWhisper*'s minimal hardware requirements and suggests the feasibility of further miniaturizing the receiver hardware.

Impact of Frequency Spacing. The frequency spacing is a crucial factor for ensuring the robustness of our FSK modulation scheme. To determine the optimal frequency spacing that can be reliably demodulated, we conduct experiments by modulating symbols using a range of frequencies between 450 Hz and 600 Hz, with a sampling rate of 19.2 kHz. Fig. 22 demonstrates that a frequency spacing larger than 100 Hz can achieve an average BER lower than 4%. This finding not only confirms the reliability of *MagWhisper*'s FSK modulation but also opens up opportunities for higher-order modulation schemes, such as 4FSK, to potentially double the transmission rate in future designs.

D. End-2-End Case Study

In this section, we present an end-to-end real-world data exfiltration attack using *MagWhisper*, exploring two practical scenarios: (a) the cross-room scenario and (b) the EM shielding scenario using a Faraday cage. Fig. 15 illustrates the experimental setup. In the cross-room scenario, the victim host and *MagWhisper* are placed in adjacent rooms, separated by a 25 cm concrete wall. In the EM shielding scenario, the victim host is securely enclosed in a Faraday cage, ensuring electromagnetic shielding, while *MagWhisper* is placed in a nearby drawer.

We exfiltrate two types of credential data: personal information text and private keys, with sizes of 512 bytes and 256 bytes, respectively. Experiments are conducted both with and without Forward Error Correction (FEC) using a 7-4 Hamming Code, at bit rates of 72 bps and 41 bps, respectively. The resulting BERs are shown in Table I. Without FEC, the BERs are 2.7% and 0.9%, for the two scenarios. With FEC, the BER drops to below 0.1%. Table II compares the recovered data against the original credentials. Without FEC, the exfiltrated personal information text contains minor errors that can be corrected using contextual knowledge. However, for data without inherent redundancy, such as private keys, FEC proves to be highly beneficial. In all FEC-enabled cases, the recovered data matches the original perfectly. These results highlight the effectiveness of *MagWhisper* in realistic, constrained environments such as through-wall and shielded scenarios, demonstrating its practicality as a powerful tool for covert data exfiltration.

E. Comparison with Existing Systems

We compare *MagWhisper* with representative covert channel systems leveraging different physical media, as summarized in Table III. BitWhisper [3] and Fansmitter [4], which utilize thermal and fan-induced acoustic artifacts, are severely constrained by their low throughput, typically below 0.25 bps. Deaf-aid [6] achieves higher throughput but relies on a gyroscope-equipped receiver and requires line-of-sight communication, making it susceptible to detection. Systems based on EM emanation or mmWave sensing, exemplified by SpiralSpy [5] and GSMem [8], support greater transmission distances. However, they are vulnerable to EM shielding. In practice, many computers are enclosed in metallic cases or placed inside Faraday cages, which block such emissions.

In contrast, *MagWhisper* overcomes these limitations by leveraging the non-radiative magnetic emanations from GPUs, which are unaffected by EM shielding. Compared with existing magnetic-based solutions using CPUs, i.e., ODINI [7] and Magneto [10], ***MagWhisper demonstrates a 13x higher bit rate and longer communication distance.*** Additionally, *MagWhisper* provides a cost-effective advantage as the sensor used in ODINI [23] is 550x more expensive than *MagWhisper*. The superior performance of *MagWhisper* can be attributed to its unique contribution in designing high-frequency FSK modulation schemes. Specifically, *MagWhisper* employs *operation-level* switching between arithmetic and memory operations to

TABLE III: Comparison between *MagWhisper* and existing data exfiltration systems.

System	Modality	Modulation	Distance	EM Shielding Resilience	NLOS Compatibility	Bit Rate
BitWhisper [3]	Thermal	ASK	0.4 m	✓		≪1 bps
Fansmitter [4]	Acoustic	ASK	1 m	✓		0.25 bps
Deaf-Aid [6]	Ultrasonic	ASK	0.2 m	✓		47 bps
SpiralSpy [5]	mmWave	ASK	8 m		✓	2 bps
GSMem [8]	EM	ASK	1 m			2 bps
ODINI [7]	Magnetic	FSK	0.2 m	✓	✓	10 bps
Magneto [10]	Magnetic	FSK	0.1 m	✓		0.2 bps
***MagWhisper* (Ours)**	**Magnetic**	**FSK**	**0.35 m (Up to 14 m in §VIII)**	✓	✓	**133 bps**

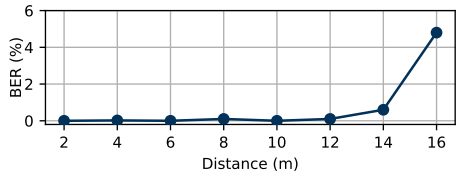

Fig. 23: Figure depicts that without changing *MagWhisper*'s transmitter, the exfiltrated data can also be received by EM probes from a distance of up to 14 meters.

generate FSK signals in the kilohertz range. In contrast, ODINI and Magneto rely on *function-level* switching, limiting their FSK modulation to just 3–7 Hz and resulting in significantly lower data rates. Moreover, in §VIII, we demonstrate that *MagWhisper*'s communication distance can be extended to 14 m. These advantages make *MagWhisper* a more effective and practical data exfiltration framework in real-world scenarios.

VIII. Long-Range Extension of *MagWhisper*

MagWhisper is a generic data exfiltration framework using GPUs, which serve as compound physical leakage sources (as discussed in §II-B). While we primarily focus on the GPU's non-radiative magnetic side channel, *MagWhisper* can also exploit high-frequency EM emanations from GPU DRAM to enable **long-range** attacks **without modifying the transmitter**. Specifically, while *MagWhisper* modulates non-radiative magnetic signals, it also *simultaneously modulates GHz-range EM emissions* from GPU DRAM [24]. This multi-band property enables an alternative receiver design using EM probes to capture long-range GHz signals. To demonstrate this feasibility, we implemented an EM receiver using a USRP B210 with a directional antenna, configured with 384 kHz bandwidth and a center frequency of 5.5 GHz. The transmitter remains unmodified, operating at a bit rate of 133 bps. As shown in Fig. 23, *MagWhisper* achieves a BER below 0.7% **at distances up to 14 meters, significantly exceeding the effective range of existing EM-based covert channels** [8], owing to the potent signal leakage from GPU hosts. Therefore, when the host is not EM shielded, attacker can launch long-range attack to further enhance stealthiness during data exfiltration.

IX. Related Work

Magnetic Side-Channels. Previous research has explored the use of magnetic leakage signals for various purposes, including inferring or eavesdropping on sensitive information like cryptography keys [25], audio [26], stylus pen writing [27], device interactions [28], [29], payment tokens [30], [31], intellectual property [32], and CPU information [33]. Notably, one study focuses on leveraging GPU magnetic signals to infer neural network architecture by identifying network layers and activation function types based on GPU synchronization points [34]. In comparison, *MagWhisper*'s unique contribution lies in the design of a robust covert channel with a high transmission capacity for conducting data exfiltration attacks.

Covert Channels for Data Exfiltration. Researchers have explored various physical artifacts for covert data exfiltration. Inertial sensors were considered, but they often require physical contact, specialized equipment, or artificial assistance [35], [36], [37]. Non-contact methods typically utilize artifacts including thermal [3], [38], light [39], [40], ultrasound [6], magnetic [10], [7], and EM-based methods [8], [41], [7], [42], as discussed in §VII-E. Compared with these methods, *MagWhisper* offers a robust covert channel with a remarkable bit rate and a low attack cost, while also enabling efficient data exfiltration under NLoS and EM-shielded scenarios. These attributes make *MagWhisper* a powerful and versatile solution for data exfiltration in various environments.

X. Conclusion

This paper presents *MagWhisper*, a novel data exfiltration attack against air-gapped hosts that leverages non-radiative magnetic emanations from GPUs. We design and implement *MagWhisper*, and conduct real-world evaluations across heterogeneous GPUs under diverse conditions. *MagWhisper* outperforms existing exfiltration techniques in bit rate and demonstrates strong robustness under NLoS and EM shielding scenarios. Furthermore, *MagWhisper* can be extended to long-distance attacks without modifying the transmitter, broadening its applicability across various threat models.

Acknowledgement

This paper is supported by the National Key R&D Program of China (Grant No.2023YFA1009500), the Fundamental Research Funds for the Central Universities (Grant No.2025110528-0), and the National Natural Science Foundation of China under grant 62372400.

REFERENCES

[1] ZDNET, "New malware makes air-gapped data center networks less bulletproof," https://www.datacenterknowledge.com/security/new-malware-makes-air-gapped-data-center-networks-less-bulletproof, 2020.

[2] OpenAI, "Reimagining secure infrastructure for advanced ai," https://openai.com/index/reimagining-secure-infrastructure-for-advanced-ai/, 2024.

[3] M. Guri, M. Monitz, Y. Mirski, and Y. Elovici, "Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations," in *IEEE 28th Computer Security Foundations Symposium, CSF*. IEEE Computer Society, 2015, pp. 276–289.

[4] M. Guri, Y. A. Solewicz, and Y. Elovici, "Fansmitter: Acoustic data exfiltration from air-gapped computers via fans noise," *Comput. Secur.*, vol. 91, p. 101721, 2020.

[5] Z. Li, B. Chen, X. Chen, H. Li, C. Xu, F. Lin, C. X. Lu, K. Ren, and W. Xu, "Spiralspy: Exploring a stealthy and practical covert channel to attack air-gapped computing devices via mmwave sensing," in *29th Annual Network and Distributed System Security Symposium (NDSS)*. The Internet Society, 2022.

[6] M. Gao, F. Lin, W. Xu, M. Nuermaimaiti, J. Han, W. Xu, and K. Ren, "Deaf-aid: mobile iot communication exploiting stealthy speaker-to-gyroscope channel," in *The 26th Annual International Conference on Mobile Computing and Networking (MobiCom)*. ACM, 2020, pp. 53:1–53:13.

[7] M. Guri, B. Zadov, and Y. Elovici, "ODINI: escaping sensitive data from faraday-caged, air-gapped computers via magnetic fields," *IEEE Trans. Inf. Forensics Secur.*, vol. 15, pp. 1190–1203, 2020.

[8] M. Guri, A. Kachlon, O. Hasson, G. Kedma, Y. Mirsky, and Y. Elovici, "Gsmem: Data exfiltration from air-gapped computers over GSM frequencies," in *24th USENIX Security Symposium, USENIX Security 15*. USENIX Association, 2015, pp. 849–864.

[9] A. Goldsmith, *Wireless communications*. Cambridge university press, 2005.

[10] M. Guri, "MAGNETO: covert channel between air-gapped systems and nearby smartphones via cpu-generated magnetic fields," *Future Gener. Comput. Syst.*, vol. 115, pp. 115–125, 2021.

[11] J. D. Jackson, *Classical Electrodynamics; 2nd Edition*. New York, NY: Wiley, 1975.

[12] Texas Instrument, "DRV425 - fully-integrated fluxgate magnetic sensor for open-loop applications," https://www.ti.com/product/DRV425, accessed: 2023-02-12, 2023.

[13] ——, "ADS1263," https://www.ti.com/lit/ds/symlink/ads1263.pdf, accessed: 2022-05-02, 2022.

[14] Raspberry Pi, "Raspberry Pi 4," https://www.raspberrypi.com/products/raspberry-pi-4-model-b/, accessed: 2022-05-02, 2022.

[15] Entrepreneur, "The rising threat of generative ai in social engineering cyber attacks," https://www.entrepreneur.com/science-technology/how-cyber-criminals-are-weaponizing-generative-ai/455896, accessed: 2023-07-30, 2023.

[16] Microsoft, "Supply chain attacks," https://learn.microsoft.com/en-us/microsoft-365/security/intelligence/supply-chain-malware, accessed: 2023-07-30, 2023.

[17] ——, "How malware can infect your PC," https://support.microsoft.com/en-us/windows/how-malware-can-infect-your-pc/-872bf025-623d-735d-1033-ea4d456fb76b, accessed: 2023-07-30, 2023.

[18] Mimecast, "Malicious insider," https://www.crowdstrike.com/cybersecurity-101/insider-threats/, accessed: 2023-07-30, 2023.

[19] TechPowerUp, "GPU-Z," https://www.techpowerup.com/gpuz/, accessed: 2022-05-02, 2022.

[20] E. O. Brigham and R. Morrow, "The fast fourier transform," *IEEE spectrum*, vol. 4, no. 12, pp. 63–70, 1967.

[21] T. Rakthanmanon, B. Campana, A. Mueen, G. Batista, B. Westover, Q. Zhu, J. Zakaria, and E. Keogh, "Searching and mining trillions of time series subsequences under dynamic time warping," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2012, pp. 262–270.

[22] NVIDIA, "Cuda toolkit," https://developer.nvidia.com/cuda-toolkit/, accessed: 2022-05-02, 2022.

[23] Newark, "Honeywell HMR2300," https://www.newark.com/honeywell-m-ps/hmr2300-d21-232/magnetometer/dp/16F5423, 2023.

[24] Z. Zhan, Z. Zhang, S. Liang, F. Yao, and X. D. Koutsoukos, "Graphics peeping unit: Exploiting EM side-channel information of gpus to eavesdrop on your neighbors," in *43rd IEEE Symposium on Security and Privacy, SP*. IEEE, 2022, pp. 1440–1457.

[25] D. Genkin, L. Pachmanov, I. Pipman, and E. Tromer, "Stealing keys from pcs using a radio: Cheap electromagnetic attacks on windowed exponentiation," in *Cryptographic Hardware and Embedded Systems (CHES)*, vol. 9293. Springer, 2015, pp. 207–228.

[26] Q. Liao, Y. Huang, Y. Huang, Y. Zhong, H. Jin, and K. Wu, "Magear: eavesdropping via audio recovery using magnetic side channel," in *The 20th Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*. ACM, 2022, pp. 371–383.

[27] Y. Liu, K. Huang, X. Song, B. Yang, and W. Gao, "Maghacker: eavesdropping on stylus pen writing via magnetic sensing from commodity mobile devices," in *18th International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, 2020, pp. 148–160.

[28] E. J. Wang, T. Lee, A. Mariakakis, M. Goel, S. Gupta, and S. N. Patel, "Magnifisense: inferring device interaction using wrist-worn passive magneto-inductive sensors," in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. ACM, 2015, pp. 15–26.

[29] R. Xiao, T. Li, S. Ramesh, J. Han, and J. Han, "Magtracer: Detecting GPU cryptojacking attacks via magnetic leakage signals," in *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking (MobiCom)*. ACM, 2023, pp. 68:1–68:15.

[30] M. Choi, S. Oh, I. Kim, and H. Kim, "Magsnoop: listening to sounds induced by magnetic field fluctuations to infer mobile payment tokens," in *MobiSys '22: The 20th Annual International Conference on Mobile Systems, Applications and Services*. ACM, 2022, pp. 409–421.

[31] D. Choi and Y. Lee, "Eavesdropping one-time tokens over magnetic secure transmission in samsung pay," in *10th USENIX Workshop on Offensive Technologies, WOOT*. USENIX Association, 2016.

[32] C. Song, F. Lin, Z. Ba, K. Ren, C. Zhou, and W. Xu, "My smartphone knows what you print: Exploring smartphone-based side-channel attacks against 3d printers," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2016, pp. 895–907.

[33] Y. Cheng, X. Ji, J. Zhang, W. Xu, and Y. Chen, "Demicpu: Device fingerprinting with magnetic signals radiated by CPU," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2019, pp. 1149–1170.

[34] H. T. Maia, C. Xiao, D. Li, E. Grinspun, and C. Zheng, "Can one hear the shape of a neural network?: Snooping the GPU via magnetic side channel," in *31st USENIX Security Symposium, USENIX Security 2022*. USENIX Association, 2022, pp. 4383–4400.

[35] N. Roy, M. Gowda, and R. R. Choudhury, "Ripple: Communicating through physical vibration," in *12th USENIX Symposium on Networked Systems Design and Implementation, (NSDI)*. USENIX Association, 2015, pp. 265–278.

[36] N. Roy and R. R. Choudhury, "Ripple II: faster communication through physical vibration," in *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, K. J. Argyraki and R. Isaacs, Eds. USENIX Association, 2016, pp. 671–684.

[37] B. Farshteindiker, N. Hasidim, A. Grosz, and Y. Oren, "How to phone home with someone else's phone: Information exfiltration using intentional sound noise on gyroscopic sensors," in *10th USENIX Workshop on Offensive Technologies (WOOT)*. USENIX Association, 2016.

[38] D. B. Bartolini, P. Miedl, and L. Thiele, "On the capacity of thermal covert channels in multicores," in *Proceedings of the Eleventh European Conference on Computer Systems (EuroSys)*. ACM, pp. 24:1–24:16.

[39] A. Maiti and M. Jadliwala, "Light ears: Information leakage via smart lights," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 3, no. 3, pp. 98:1–98:27, 2019.

[40] M. Guri, B. Zadov, and Y. Elovici, "Led-it-go: Leaking (A lot of) data from air-gapped computers via the (small) hard drive LED," in *Detection of Intrusions and Malware, and Vulnerability Assessment - 14th International Conference, DIMVA*, vol. 10327. Springer, 2017, pp. 161–184.

[41] M. Guri, M. Monitz, and Y. Elovici, "Usbee: Air-gap covert-channel via electromagnetic emission from USB," in *14th Annual Conference on Privacy, Security and Trust (PST)*. IEEE, 2016, pp. 264–268.

[42] N. Sehatbakhsh, B. B. Yilmaz, A. G. Zajic, and M. Prvulovic, "A new side-channel vulnerability on modern computers by exploiting electromagnetic emanations from the power management unit," in *IEEE International Symposium on High Performance Computer Architecture, (HPCA)*. IEEE, 2020, pp. 123–138.