

VReaves: Eavesdropping on Virtual Reality App Identity and Activity via Electromagnetic Side Channels

Wei Sun¹, Minghong Fang², Mengyuan Li³

redsunwit@gmail.com, minghong.fang@louisville.edu, mengyuanli@usc.edu

¹Wichita State University, ²University of Louisville, ³University of Southern California

Abstract

Virtual reality (VR) has recently proliferated significantly, consisting of headsets or head-mounted displays (HMDs) and hand controllers for an embodied and immersive experience. The VR device is usually embedded with different kinds of IoT sensors, such as cameras, microphones, communication sensors, etc. However, VR security has not been scrutinized from a physical hardware point of view, especially electromagnetic emanations (EM) that are automatically and unintentionally emitted from the VR headset. This paper presents VReaves, a system that can eavesdrop on the electromagnetic emanation side channel of a VR headset for VR app identification and activity recognition. To do so, we first characterize the electromagnetic emanations from the embedded IoT sensors (e.g., cameras and microphones) in the VR headset through a signal processing pipeline and further propose machine learning models to identify the VR app and recognize the VR app activities. Our experimental evaluation with commercial off-the-shelf VR devices demonstrates the efficiency of VR app identification and activity recognition via electromagnetic emanation side channel.

Keywords

Virtual Reality, Electromagnetic Emanations, App Identification and Activity Recognition

1 Introduction

Virtual reality (VR) systems have become omnipresent in our daily lives, usually consisting of a VR headset or head-mounted display (HMD) and two hand controllers. We can use these VR systems for an embodied and immersive user experience, such as playing video games, watching videos, online training, collaboration, etc. Among all of these, virtual reality gaming [39, 55] has shown growing interest due to the proliferation of mobile computing and human-computer interaction. As such, there are different kinds of commercial off-the-shelf VR systems being developed on the market, such as Meta Quest [33], Sony PlayStation VR2 [49], HTC Vive [19], Apple Vision Pro [2], etc. As a mobile device, the VR headset can run different VR apps, which can reveal the VR user's private information (e.g., personality [21, 41] and behavior biometrics [40]). These VR platforms are embedded with different types of IoT sensors [3], such as cameras and microphones for an embodied and immersive user experience, which are vulnerable to different side channel-based privacy attacks [11, 47, 48, 66].

However, the existing side channel-based attacks on VR systems mainly focus on acoustics [30], VR user behaviors (e.g., head movements [36, 66] or hand gestures [15]), RF side channel [1], motion sensors [35], camera-captured videos/images (e.g., user avatar) [58],

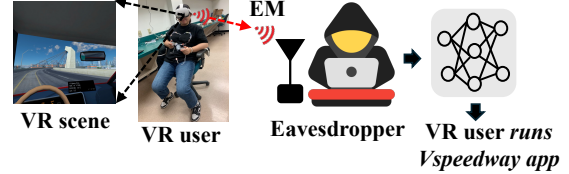


Figure 1: VR headset leaks emanations during operation, which can be sniffed by the eavesdropper to infer the VR app's identities (e.g., Vspeedway app and Aim app) and activities (e.g., app configuration and running).

network traffic [51], etc., which leave the electromagnetic side-channel unexploited yet for VR platforms. Moreover, all these attacks either require malware or a proactive surveillance implant in the VR user's environment to obtain private user-related (e.g., motion-related) information. Therefore, it is essential to exploit the automatic and unintentional electromagnetic side channel to scrutinize the VR platform.

To this end, this paper presents VReaves, a system that can exploit the electromagnetic emanations (EM) automatically and unintentionally generated by the IoT devices (e.g., camera, microphone) embedded in the VR headset to eavesdrop on the VR headset or head mount display (HMD) as shown in Fig. 1, which can reveal important private information (e.g., personalities) related to the VR users through VR app identification and activity recognition. Therefore, this EM side channel can pose a great privacy threat to the VR users. For example, people can make a profit from understanding VR app identities and activities by leveraging this information to make targeted app recommendations and even reveal the user's personality and daily living habits [40, 41].

To have functional and practical EM-based eavesdropping on the VR device, the adversary can use a wireless receiver (e.g., software-defined radio) to eavesdrop on the emanations emitted from the VR user's headset for VR app information analysis. However, there are three great challenges. First, even though the emanation as side channel information is widely exploited for hidden camera or microphone detection [52, 64, 67], camera image reconstruction [28], and screen information reading [26], the prior works do not reveal the relationship between the emanations from VR headset and the VR app identities and activities. Different from the prior work of characterizing the single emanation source (e.g., camera or microphone), the VR headset consists of different kinds of emanation sources. As a result, the emanations from these sources are interleaved with each other, which are difficult to extract and characterize using the techniques proposed in the prior works. Therefore, we not only need to accurately extract the frequency-domain emanations but

also properly characterize the emanations for the VR app identification and activity recognition. To do so, we propose to use fine-tuned machine learning models to characterize the relationships between the interleaved emanations and VR app information.

Second, the emanations are unintentionally and automatically emitted from the VR headset, which can be weak in strength and easily interfere with the ambient wireless signals. This is because the emanations are amplitude-modulated clock signals in the time domain that are spreading across a wide frequency band. Therefore, it is important to boost the emanation strength and suppress the ambient wireless interference. To address this challenge, we propose a signal-processing pipeline, including noise floor smoothing, interference suppression, and emanation strength enhancement to characterize the emanation spectrum accurately. Then, we design a multi-frequency machine learning model to characterize the frequency-domain emanations for VR app identification.

At last, it is difficult to infer the VR app identities and activities based on the same frequency-domain emanations simultaneously. This is because the frequency-domain emanations exhibit the same pattern for a specific VR app, which cannot be used to discriminate the fine-grained app activities. Therefore, we further explore the over-time frequency-domain emanations to characterize the VR app's activities. Specifically, we propose to use short-time frequency transform (STFT) to derive the spectrogram of the emanations and further regard it as an image for the VR app activity recognition using a multi-spectrogram machine learning model.

To demonstrate the efficiency of the emanation-based VR app identification and activity recognition, we built a prototype with the software-defined radio (i.e., USRP N210) instrumented with the LP1401 directional antenna for emanation eavesdropping. Our extensive system performance evaluation with commercial off-the-shelf VR platforms (e.g., Meta Quest 3 and HTC VIVE XR Elite) achieves an accuracy of 99% in the VR app identification and an accuracy of 99% in the VR app activity recognition under various settings including distance, orientation, hardware and software configurations, etc. We summarize the main contribution of our system design in the following.

- To the best of our knowledge, this is the first system that exploits the emanations from the VR headset for VR app identification and activity recognition.
- We propose a signal processing pipeline to smooth the noise floor across a wide frequency band, suppress ambient wireless interference, and boost the emanation strength for accurate emanation extraction and characterization.
- To reveal the relationship between the emanations and VR app identities and activities, we propose to characterize the frequency-domain emanations and over-time frequency-domain emanations through averaging FFT and STFT, using the fine-tuned pre-trained machine learning models.
- Our experimental evaluations with system-level tests, case study, and microbenchmarks demonstrate the efficiency of the VR app identification and activity recognition based on the characterized emanation leakage.

This paper marks an important step in establishing the relationship between the emanations and computational activities in the VR device, which can propel the field forward on VR device scrutiny.

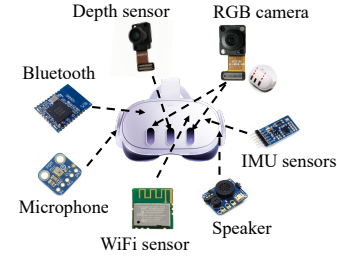


Figure 2: The sensors embedded in the Meta Quest 3 (or HTC Vive XR Elite) include cameras, speakers, microphones, WiFi, and Bluetooth sensors that can be the potential emanation sources.

Since emanations are modulated by the computational activities, we can leverage these emanations to infer the videos in VR and further reconstruct the VR scenes displayed in the headset.

2 Background

2.1 VR Devices and Their Emanations

The virtual reality platforms (e.g., Meta Quest, HTC Vive XR Elite, etc.) usually consist of a headset (or head mount display) and two hand controllers. For example, as shown in Fig. 2, the Meta Quest 3 headset is embedded with different types of sensors for perception, display, and detection of physical scenes. These embedded sensors can automatically and unintentionally emit electromagnetic emanations, which can be used as a side channel to steal human private information. We primarily summarize five types of emanation sources from the VR headset in the following.

- **Camera’s emanations.** The camera sensors are connected to the CPU, GPU, or image processing unit for raw pixel data transmission through High-speed Serial Pixel Interface, Digital Video Port, Low-voltage Differential Signaling, or MIPI Camera Serial Interface 2, which can leak the emanations [28, 52] that can indicate the computation activities of the camera sensor as shown in Fig. 3(1).
- **Display’s emanations.** The display in the headset could leak emanations [26, 28], consisting of the graphical computing unit (GCU) and screen (e.g., OLED or LED displays). The emanations are emitted from the data interface connecting the GCU and screen as shown in Fig. 3(2).
- **Microphone and speaker’s emanations.** The microphones and speakers can be the emanation sources [7, 67], which usually include the ADC or DAC controlled by the clock signals for synchronization that can leak emanations as shown in Fig. 3(3).
- **RF radio’s emanations.** The wireless communication radios (e.g., WiFi and Bluetooth radios) can emit emanations [4, 6] through the oscillators or mixers as shown in Fig. 3(4).
- **Memory’s emanations.** The dynamic random-access memory (DRAM) can introduce the emanations [45, 46, 62] in modern electronic devices due to memory access operation of the CPU as shown in Fig. 3(5).

The emanations are emitted through the circuits and data/signal pipelines on these sensors that can be unintentionally regarded

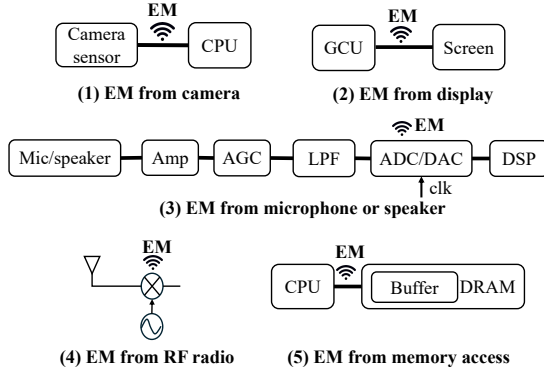


Figure 3: Emanation source in VR headset includes (1) camera sensors, (2) screens or monitors, (3) microphones or speakers, (4) wireless communication radios (e.g., WiFi and Bluetooth sensors), and (5) dynamic random-access memory.

as radio-frequency antennas. Next, we illustrate the primer and properties of the emanations.

2.2 Primer on Electromagnetic Emanations

Physical principle of EM. As shown in Fig. 4, the electromagnetic emanations (or emanations) are amplitude-modulated clock signals, as the clock signals are coupled with the computational activities (e.g., switching behaviors) through hardware components such as the capacitor, resistor, or diodes. The clock signals can be expressed as follows:

$$s_{clk}(t) = \cos(2\pi f_0 t + \frac{\Delta f}{f_m} \sin(2\pi f_m t)) \quad (1)$$

where f_0 is the clock frequency, f_m is modulating frequency, and Δf is the peak frequency deviation. As a result, the frequency-domain clock signals can be derived as follows:

$$\|F(f)\| = \left\| \sum_n J_n \left(\frac{\Delta f}{f_m} \right) \times (\delta(f - f_0 + n f_m) - \delta(f - f_0 - n f_m)) \right\| \quad (2)$$

where $J_n(\cdot)$ is the Bessel function and $\delta(\cdot)$ is the Dirac delta function. For the sake of simplicity, we can rewrite the spectrum expression in the above as follows:

$$s_{clk}(t) = \sum_{n=1}^N A_{clk}(n) \sin(2\pi f_{clk} t) \quad (3)$$

where $A_{em}(n)$ represents the amplitude of the spikes and $f_{clk} = f_0 - n f_m$, or $f_{clk} = f_0 + n f_m$. Now, let's consider the computational activity, such as a series of periodic memory accesses or switching behaviors, which can introduce the square waves at the clock edges. This square wave can be expressed with the Fourier transform as follows:

$$s_{sq}(t) = \frac{2A_{sq}}{\pi} \sum_m \frac{\cos((2m-1)2\pi f_{sq} t)}{2m-1} \quad (4)$$

where f_{sq} and A_{sq} denote the frequency and amplitude of the square wave introduced by the computational activity. As a result, this computational activity introduces the amplitude modulation to the

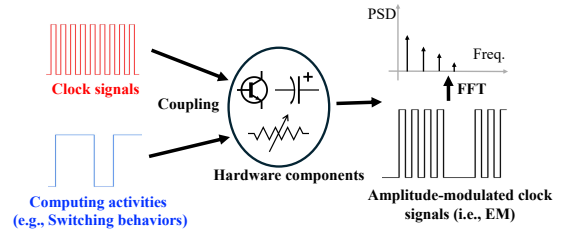


Figure 4: Electromagnetic emanations are amplitude-modulated clock signals, introduced by the clock signals coupling with the computational activities.

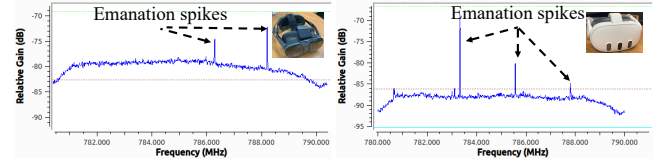


Figure 5: Frequency-domain representation of emanations from HTC VIVE XR Elite and Meta Quest 3.

clock signals, which is equivalent to frequency mixing. Therefore, the emanations are produced as:

$$s_{em}(t) = s_{clk}(t) s_{sq}(t) = \sum_n \sum_m \frac{A_{clk}(n) A_{sq}}{(2m-1)\pi} s_p(n, m, t) \quad (5)$$

where $s_p(n, m, t) = \sin(2\pi(f_{clk} + (2m-1)f_{sq})t) + \sin(2\pi(f_{clk} - (2m-1)f_{sq})t)$ indicates clock-introduced spikes are modulated by the computational activity-introduced spikes.

EM characterizations. The emanations are clock signals that are modulated in amplitude by the computational activities or switching behaviors on the hardware devices. The clock signals are coupled with the switching behaviors or computational activities through different hardware components such as capacitors, resistors, and diodes. As a result, the emanations in the time domain exhibit the squared waves, where the periodicity is the clock frequency. In the frequency domain, the emanations exhibiting the harmonics can spread across a wide frequency band with multiples of emanation spikes, consisting of a fundamental harmonic and multiples of harmonics as shown in Fig. 5. The distance between the adjacent harmonics indicates the clock frequency.

EM security and privacy. As we can see, the spreading emanation spectrum consists of a fundamental frequency spike and a series of multiple frequency spikes, which are determined by the computational activities at the emanation source (e.g., VR headset). These emanations of the VR headset can carry important information about the VR users, which can be carefully characterized through signal processing techniques (e.g., FFT or STFT). Then, we can leverage the machine learning models to reveal the hidden privacy information from the emanations, resulting in a great privacy threat. However, since the energy of the emanations is spread across a wide spectrum, it is difficult to detect all the emanation harmonics emitted from different emanation sources like the prior works [20, 52, 57, 59]. Moreover, the ambient wireless communication signals can interfere with the emanations, which need to

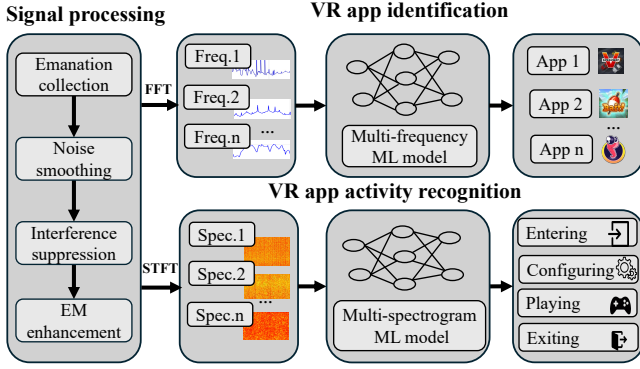


Figure 6: The workflow of VReaves consists of the emanation signal processing module, VR app identification module, and VR app activity recognition module.

be suppressed for accurate harmonic extraction. The interleaved emanations from the different sensors in the VR headset make the privacy attack more challenging.

3 Threat Model

Attack settings. Our attack leverages the wireless signal receiver (e.g., software-defined radio USRP N210, PlutoSDR, LimeSDR, etc.) to passively sniff the electromagnetic emanations emitted from the VR headset. The VR user can wear the VR headset or the head-mounted display (HMD) to enjoy the embodied and immersive experience. We assume that the attacker can deploy a radio-frequency sniffer close to the victim for emanation eavesdropping. For example, the VR user plays the VR games in a public space (e.g., cafeteria, transportation) while the attacker is eavesdropping next to the victim [17]. Similar attack scenarios were validated in emanation-based attacks [7, 8, 17, 30] for other purposes.

Adversarial model. We make the following assumptions about the adversary.

- The adversary does not have physical or remote access to the victim’s VR platform. However, the adversary can deploy the eavesdropping radio to sniff the electromagnetic emanations emitted from the VR headset.
- The adversary only performs passive emanation eavesdropping using the RF radio instrumented with the directional antenna, which connects with the laptop to run the algorithms for VR app identification and activity recognition.
- The goal of the attacker is to predict the VR app’s identities (e.g., Aim app and Vspeedway app) and activities (e.g., entering or configuring the VR app), which can be further abused for app recommendation and the VR user’s personality inference.
- Similar to the attack proposed in [37], we assume the adversary mainly focuses on a set of VR apps and app activities. The inference of the VR user’s personality and biometric behavior using eavesdropped VR app identities and activities is beyond the scope of this paper.

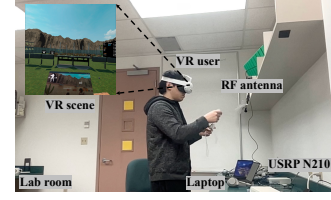


Figure 7: Experimental setup in the lab room, where the VR user wears the Meta Quest 3, and the USRP N210 connects with the directional antenna for emanation sniffing.

4 System Overview

Fig. 6 shows the workflow of our eavesdropping system consisting of the signal processing module, app identification module, and app activity recognition module.

- **Signal processing module.** After the emanation signals are eavesdropped across a wide frequency band, we first smooth the noise floor with a mean filter to remove the frequency-dependent noise variation across the frequency bands. Then, we suppress the ambient interferences, including the ambient emanations and wireless communication signals, through spectrum subtraction. We further boost the emanation strength through non-coherent averaging of the eavesdropped signals over time.
- **VR app identification module.** To identify the VR app, we design a fine-tuned pre-trained multi-frequency ML model to infer the VR app identities based on frequency-domain characteristics of the emanations using FFT.
- **VR app activity recognition module.** To recognize the VR app activities, we design a fine-tuned pre-trained multi-spectrogram ML model to infer the app activities based on the spectrogram of emanations derived with STFT that can characterize over-time properties of the frequency-domain emanations.

In what follows, we first present the emanation signal processing, including the experimental evaluation-based noise floor smoothing, ambient wireless interference suppression, and emanation strength enhancement. Then, we establish the relationship between the characterization of emanations and the app identities or activities through machine learning models.

5 System Design

5.1 Emanation Signal Processing

The emanations from the VR headset exhibit amplitude-modulated clock signals in the time domain and spread across a wide frequency band in the frequency domain. Our goal is to accurately extract and characterize the frequency-domain emanation signals for VR app identification and activity recognition. Our attack design options are demonstrated through the experimental measurements.

Experimental setup. To do so, we illustrate the experimental setup shown in Fig. 7. The VR user is wearing the headset of Meta Quest 3 or HTC VIVE XR Elite for an embodied and immersive experience. At the same time, we use the USRP N210 instrumented with the directional antenna and UBX daughterboard to sniff emanations. The USRP N210 connects to the laptop of the Lenovo ThinkPad for

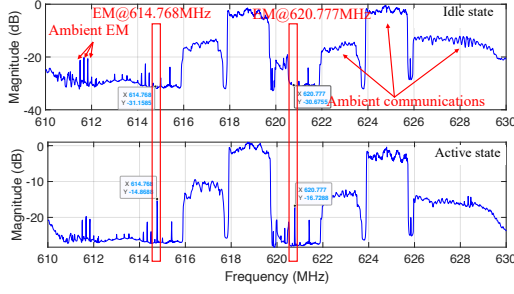


Figure 8: FFT of the eavesdropped wireless signals during active and idle states of VR application usage, showing the emanations from the VR headset, ambient emanations, and wireless communication signals.

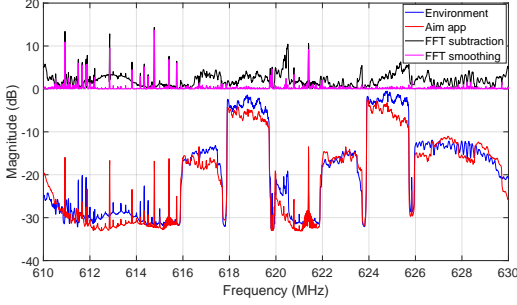


Figure 9: FFT of the emanations in the physical wireless environment during the active and idle states of application usage, and the FFT of the emanations after ambient wireless signals subtraction and noise floor smoothing.

IQ sample processing. More details about our attack system implementation can be found in Section 6. From this experimental setup, we explore the impact of noise and ambient wireless interferences on the emanations. Furthermore, we notice that the emanations should be strengthened for detection.

Noise floor smoothing. Since the frequency-domain emanations are spreading across a wide frequency band, the varying noise floor across the spectrum needs to be smoothed for accurate emanation characterization. As such, we need to scan a wide spectrum for emanation spike characterization. To do so, we first scan the frequency band below 1 GHz, using USRP N210 as a receiver instrumented with the LP0410 PCB antenna [42] with a sampling rate of 25 MHz and a bandwidth of 10 MHz. Then, we concatenate all the 10 MHz frequency bands for emanation characterization. We apply the movmedian filter [32] for noise removal. We further smooth the noise floor across a wide frequency band. After the noise removal, the noise floor across the wide frequency band becomes flat and smooth, which can advance the frequency-domain emanation characterization.

Ambient wireless interference suppression. As shown in Fig. 8, the sniffed wireless signals at the eavesdropper include the ambient emanations from the electronic devices (e.g., monitors), ambient wireless communication signals (e.g., cellular signals), and emanations from the VR headset. As we can see, the spectrum is primarily

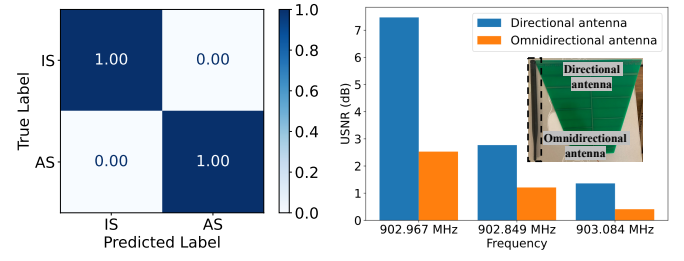


Figure 10: Confusion matrix of active state (AS) and idle state (IS) classification for VR headset.

Figure 11: Impact of the different antenna gains on USNR of the eavesdropped emanation.

dominated by ambient wireless communication signals, which need to be mitigated for accurate VR headset-introduced emanation characterization. To this end, we propose to eliminate these ambient artifacts through subtraction. Specifically, we first characterize the ambient wireless communication environment. We assume that these ambient artifacts do not change abruptly, which has been experimentally demonstrated in the prior work [52]. As such, we can eliminate these artifacts by using a sliding window-based spectrum subtraction over the sniffed wireless signals. To do so, the adversary can detect the VR app-running state through the variation of the emanations introduced by the VR headset, as the emanations are mainly affected by the computational activities on the VR headset. As a demonstration, we leverage the emanation source state detection technique from IoTProsector [53] to sense the VR app state, which can be used to guide the adversary for VR app-related emanation sensing. Fig. 10 shows the confusion matrix for the VR app state detection, where the binary classification achieves 100% accuracy. Hence, we accurately extract the VR app-related emanations for app identification and activity recognition.

Fig. 9 shows the efficiency of noise floor smoothing with the noise removal filter and interference suppression through subtraction. We first show the FFT of the received wireless signals indicated by the blue line in the figure when the VR headset is in the idle state. Then, we show the FFT of the received wireless signals indicated by the red line in the figure when the VR headset is in the active state. As we can see, most of the spectrum (e.g., the ambient wireless communication signals) is overlapped due to the quasi-stable wireless environment, while the emanation spikes are outstanding when the VR headset is in the active state. After we do subtraction between them, as indicated in the black line in the figure, the ambient wireless communication signals are suppressed. To further remove the variation of the noise floor, we apply the movmedian filter [32] on the subtracted spectrum, as indicated in the pink line. As such, the noise floor becomes flat and smooth across the frequency bands. Moreover, we can see the outstanding emanations spikes in the FFT result, which can be leveraged for app identification and activity recognition.

EM strength enhancement. Since emanations are unintentionally emitted from the VR headset, which is naturally weak in strength. The weak emanation signals restrict the emanation detection and characterization. To enhance the emanation strength received at

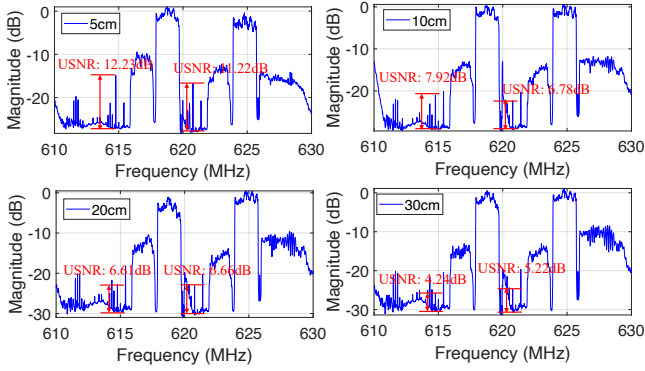


Figure 12: USNR of the emanations from the VR headset over different distances.

the eavesdropper, the straightforward idea is to use a high-gain directional antenna or amplifier. To demonstrate this, we compare the unintentional signal-to-noise ratio (USNR) [8] of the eavesdropped emanations when we use an omnidirectional antenna (i.e., VERT900 with 3dBi gain) and a directional antenna (i.e., LP0410 with 6dBi gain) instrumented on the USRP N210. Fig. 11 shows the USNR of emanation spikes using antennas with different gains. As we can see, the directional antenna with more power gains can eavesdrop on the emanations with higher strength than the omnidirectional antenna. This indicates that we can use a high-gain directional antenna for emanation eavesdropping. We can also use a power amplifier at the eavesdropper to enhance the sniffed emanation strength. Fig. 12 shows USNR over different distances between the attacker and VR user. As we can see, as the distance between the attacker and the VR headset becomes larger, the USNR becomes smaller due to the path loss. Therefore, it is important to boost the emanation signal strength for long-range emanation detection and characterization.

To further enhance the emanation strength and even bring up the emanations below the noise floor, we propose to boost the emanation signal strength by exploiting the emanation signal characteristics. Specifically, since the emanations are amplitude-modulated clock signals, they are represented as square waves in the time domain. However, the noise does not exhibit any specific pattern. So, we can use the non-coherent averaging to boost the emanation strength by taking the average of the over-time emanation signals, while the noise can be averaged out. As such, the emanations become outstanding in the spectrum.

Fig. 13 and Fig. 14 present the averaged FFT of the emanations emitted from the VR headset when we run the VR app. As we can see, when we do the averaging FFT on the eavesdropped emanations over 0.2s, the USNR of the emanation spike becomes larger in comparison to the averaging FFT on the eavesdropped emanations over 0.1s. Specifically, the emanation spike at the frequency of 614.768 MHz exhibits a USNR of 14.7874 dB when we use averaging FFT on the emanations over 0.2s, while it is 14.4809 dB when we use averaging FFT on the emanations over 0.1s. As such, we have the USNR gain of 0.2 ($= \frac{14.7874 - 14.4809}{14.4809}$) dB per second. So, averaging over a longer time duration could boost the emanation

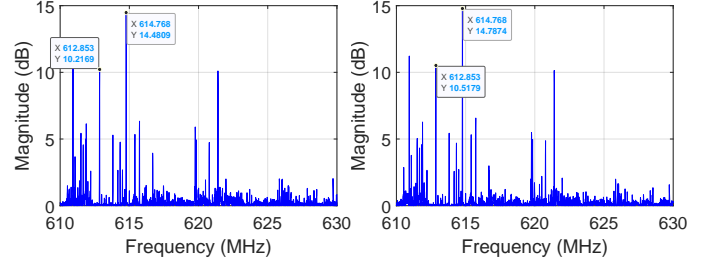


Figure 13: Averaging FFT on the emanations over 0.1s.

Figure 14: Averaging FFT on the emanations over 0.2s.

signal strength. For example, averaging over 10s could boost the emanation strength by 2 dB. This is because the emanations are square waves, while the noise does not show any specific pattern. As such, averaging the FFT could improve the emanation's USNR.

5.2 Machine Learning-based App Identities and Activities Inference

To identify the VR apps and infer their activities, we propose to use machine learning based on the characterized emanations. Specifically, we can leverage the FFT of the emanations to infer the app identities.

App identification. To demonstrate the feasibility of using FFT to discriminate different VR apps, we conduct an experiment to show the FFT results when we run different VR apps. As shown in Fig. 15, the frequency-domain emanations are different when we run different VR apps. This is because the emanations are amplitude-modulated clock signals, which are mainly affected by the computational activities on the VR headset. As a result, the frequency-domain emanations can carry VR app information, which can be used to infer VR app identities.

To accurately characterize the frequency-domain emanations, we design a neural network to classify the VR apps. Our multi-frequency neural network should characterize over-frequency emanation spikes. To this end, we propose to leverage the ResNet [54] with a fine-tuned convolutional layer that can adapt to the input of FFT across multiple frequency bands and further infer the VR app identities. This is because the pre-trained ResNet has been demonstrated to be efficient in characterizing the wireless environment. As such, we can infer the VR app identity accurately. Specifically, the input of the ResNet is the frequency-domain spikes, which have already eliminated the effect of the physical wireless environment through subtraction. The output can be used to classify app identities.

App activity recognition. To infer the app activities, can we still use the FFT of emanations to discriminate the VR app's activities? To answer this question, we conduct an experiment to show the FFT results when we configure and run a VR app. As shown in Fig. 16, the emanation spikes overlap with each other. This is because the same app is configured or operated on the VR headset. Therefore, it is not possible to differentiate the VR app's activities based on the frequency-domain emanations alone. However, we find that different activities running on the VR headset can introduce over-time

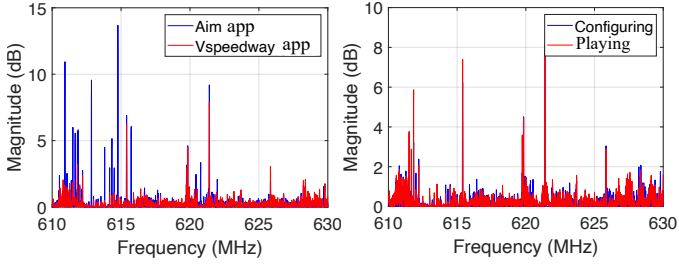


Figure 15: FFT of emanations from VR headset running different VR apps.

Figure 16: FFT of emanations from VR headset configuring or running the VR app.

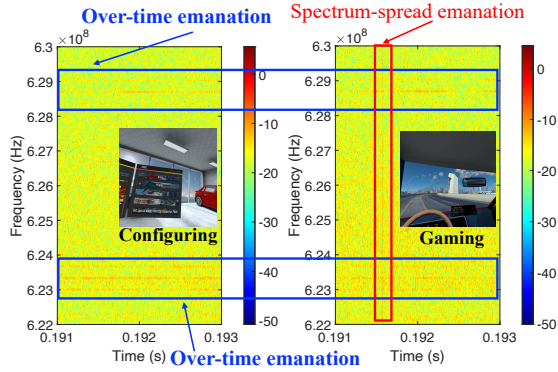


Figure 17: Spectrogram of emanations when we configure the app (left figure) and run the app (right figure). The blue rectangles show the different over-time emanation patterns for gaming and configuring. The red rectangle shows the spread-spectrum emanation only when the VR user runs the app.

characteristics to the frequency-domain emanations. To characterize the over-time properties of the emanations, we exploit the spectrogram that can show the over-time and over-frequency features of the emanations. As we can see in Fig. 17, the left and right spectrograms present the emanations in the time and frequency domain when we run and configure a VR app. As we can see, the red rectangle shows the spectrum-spread emanations across the frequency, which are presented only when the VR user runs the app. Moreover, the blue rectangles show the over-time emanations for the in-running and in-configuration VR app, which exhibit different over-time patterns. This is because the emanations are determined by the computational activities on the VR headset, resulting in differentiated over-time emanations. Therefore, we can leverage the spectrogram with time and frequency-domain features of the emanations derived with the short-time Fourier transform (STFT) to infer the VR app activities. Specifically, we leverage the multi-spectrogram machine learning model using ResNet architecture with a fine-tuned convolutional layer that can take the spectrogram as input to recognize the VR app activity, including entering, configuring, running, and exiting.



Figure 18: Experimental devices used for system performance evaluation.

Since we leverage frequency-domain emanations and the over-time frequency-domain emanations to infer app identity and activity, we cannot simply use one machine learning model for these two tasks through multi-task learning. Moreover, the app activity recognition is based on the over-time frequency characterization of the emanations. To this end, after the attacker eavesdrops on the IQ samples, the frequency-domain emanations are used to predict the app identities, and the over-time frequency-domain emanations are used to predict the app activities. These two tasks have been performed in parallel since they are independent of each other, which can potentially accelerate the attack process. Next, we illustrate the implementation and evaluation details.

6 Implementation and Evaluation

Hardware and software. To eavesdrop on the emanations from the VR headset (e.g., Meta Quest [33] or HTC VIVE XR Elite [19]), as shown in Fig. 18, we use a UBX40 daughterboard-enabled USRP N210 [43] as the software-defined radio (SDR) instrumented with the directional antenna LP0410 [42], which is also compared with the omnidirectional antenna VERT900 [44] on USNR [8] measurements. The USRP N210 connects to the ThinkPad X1 Carbon Gen11 laptop with an Intel i7 CPU running Ubuntu 20.04 OS. The sniffed IQ samples are streamed to this laptop for further analysis. Specifically, USRP N210 streams the sniffed IQ samples to the laptop, which runs our signal processing algorithms in MATLAB for emanation characterization in the frequency domain. Our fine-tuned machine learning models are well-trained and implemented with Pytorch on the server with an NVIDIA A6000 Ada GPU for app identification and activity recognition. Our Pytorch code consists of FFT and Hamming window-based STFT on the IQ samples for app identification and activity recognition, respectively. We use a fine-tuned pre-trained ResNet18 as our foundational neural network model with a fine-tuned convolutional layer and output layer.

Experimental settings. Our experimental setup consists of the VR user (i.e., victim) wearing the VR headset and an attacker eavesdropping on emanations. We consider a list of fifteen VR apps (i.e., Aim, Bait, Epic, Slupies, Vspeedway, Beast, Duck, Stable, Master, Tennis, Cosmicflow, Openbrush, Hyperdash, Maestro, and Conjure cards apps) and four app-specific user activities (i.e., entering, configuring, running, and exiting operation) that are also considered in prior work [37]. The VR users can run apps in the break room and lab room of the departmental building with rich ambient wireless signals, while the attacker eavesdrops on the emanations from the VR headset at any random locations that are 1 to 2 meters away from the VR user, as shown in Fig. 19. Our study has received IRB

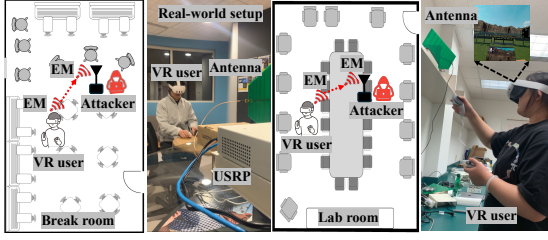


Figure 19: Experimental settings in the break room and lab room of the departmental building for system performance evaluation.

approval from our institution. We scan the spectrum below 1 GHz with a sampling rate of 25 MHz and a bandwidth of 10 MHz. We experimentally notice that the emanations are mainly emitted from the frequency band between 580 MHz and 630 MHz. As such, we mainly eavesdrop on this frequency band. Eavesdropping on more frequency bands could potentially improve the performance of emanation-based app identification and activity recognition. We explore this impact in our experimental results section. We collect 7.5G of emanation IQ samples for app identification and 42.5G for app activity recognition. By default, we report the system performance evaluation with Meta Quest 3, and the collected time-domain emanations can be divided into chunks with 500K IQ samples for FFT and STFT. Then, we split the data set into a training set with a size of 70%, a validation set with a size of 15%, and a test set with a size of 15%. The best well-trained model is used for prediction.

Evaluation metrics. To evaluate the performance of the end-to-end system, we report the confusion matrix and accuracy for VR app identification and activity recognition. We evaluate different factors that can affect system performance, such as the impact of eavesdropper-victim distance, the emanation time duration, the number of frequency bands, and different VR headsets. We also use USNR [8] as a metric to show the emanation strength. To demonstrate the efficiency of our proposed fine-tuned ResNet, we also evaluate the performance of other deep neural network models (e.g., LSTM and transformers) on the VR app identification and activity recognition. We also propose a countermeasure for this eavesdropping attack and further show its potential through simulated emanations.

7 Experimental Results

7.1 System-Level Evaluation

Method. To evaluate the end-to-end system performance, we report the performance of the VR app identification and activity recognition. Specifically, we report the system performance on app identification and activity recognition through the accuracy and confusion matrix. Our emanation measurements are sniffed in the beakroom and lab room of the departmental building, as described in the implementation and evaluation section. The fine-tuned ResNet model is well-trained on the training dataset, and the best model is selected based on the validation set for prediction on the test set. The final performance is reported based on the test

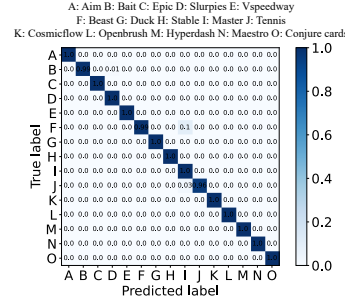


Figure 20: Confusion matrix of app identification.

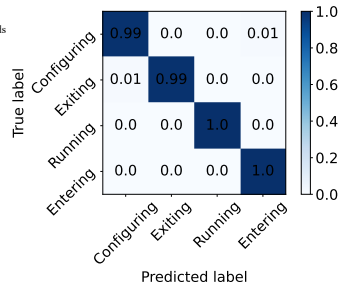


Figure 21: Confusion matrix of app activity recognition.

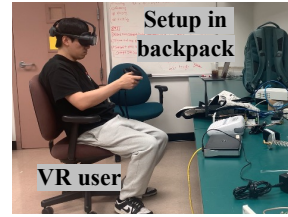


Figure 22: Concealed experimental setup in the case study.

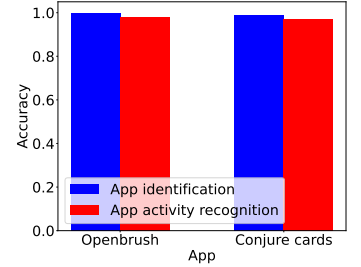


Figure 23: Performance of VReaves in the case study.

dataset. By default, we use emanations spreading across five bands for the evaluation.

Result. Fig. 20 and Fig. 21 show the confusion matrix of app identification and activity recognition. As we can see, these experimental results show the efficiency of our system in eavesdropping on VR app identities and activities via emanations. The high accuracy of the system performance is not due to the model being over-fitted, as we employ strategies (e.g., separate test and validation datasets, training adjustments, and fine-tuning pre-trained ResNet) to avoid it. Moreover, we use a fine-tuned ResNet model that can accurately characterize our signal processing-extracted emanations. Even though we use different fine-tuned ResNet models for app identification and app activity recognition, they are different from the input (i.e., FFT v.s. STFT results). As such, our attack can be conducted independently. This indicates the efficiency of the attack in identifying VR apps and recognizing app activities.

7.2 Case Study

To demonstrate the effectiveness of a stealthy attack using our proposed VReaves system, we conduct a case study. Similar to the attack scenarios reported in [17], in a typical public space, such as a cafe, a library, or a railway station, when VR users use VR devices to play games or chat online, the attacker can hide the eavesdropping setup in the backpack to closely sniff the emanations from the VR headset for app identification and activity recognition.

Method. As an illustration, we can consider the scenario as shown in Fig. 22, where the attack setup is hidden in the backpack to ensure the stealthiness and further eavesdrop on the emanation

Table 1: Performance comparison over different ML models.

ML models	Accuracy	
	App identification	Activity recognition
LSTM	0.73	0.71
Transformer	0.76	0.75
Our model	0.99	0.99

measurements from the VR headset and the VR user is running the VR app. The distance between the VR headset and the backpack is within 1 meter. This concealed setup is practical, which can be used in the public space (e.g., a cafe or a library) for eavesdropping similar to the attack scenario proposed in [17]. We use the well-trained model in the subsection 7.1 to predict the VR app identities and activities based on the emanation measurements collected with this concealed attack setup.

Result. Fig. 23 shows the accuracy of app identification and activity recognition in the concealed setup. As we can see, the accuracy of app identification and activity recognition is around 0.99, which indicates the superiority of VReaves in the concealed or hidden setup. This is because VReaves is powered by the emanation enhancement techniques proposed in Section 5 and can be deployed closely to the VR users.

7.3 Microbenchmarks

7.3.1 Impact of VR headsets. To see the impact of the VR headset on the system performance, we evaluate VR app identification and activity recognition across different VR headsets.

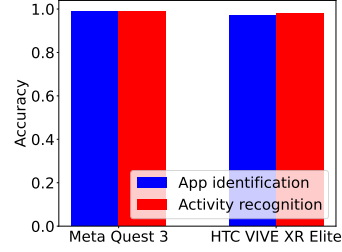
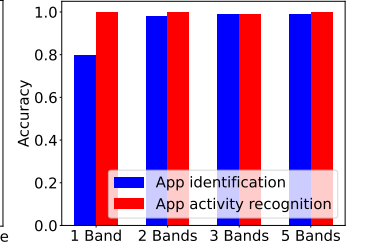
Method. To do so, we mainly use Meta Quest 3 and HTC Vive XR Elite to evaluate the impact of different VR headsets on VR app identification and activity recognition. Since the different brands of VR headsets do not share the same VR apps, we train the model on the emanation measurements collected and report the performance on the emanation measurements collected from different VR headsets separately.

Result. Fig. 24 shows the accuracy of app identification and activity recognition with Meta Quest 3 and HTC VIVE XR Elite. As we can see, the accuracy of app identification and activity recognition over different VR headsets exhibits almost the same high values, as our subtraction method could not only suppress the ambient wireless interference but also eliminate the hardware-dependent artifacts. This demonstrates the efficiency of our proposed attack system over different VR headsets.

7.3.2 Impact of ML models. Since the emanation measurements exhibit a specific pattern, we need to exploit the impact of the different machine learning models on characterizing the VR emanations.

Method. To demonstrate the efficiency of using fine-tuned ResNet for VR app identification and activity recognition, we compare it with the customized LSTM and transformer models that have already shown strong capability of characterizing the wireless spectrum.

Result. As shown in Table 1, our fine-tuned ResNet model exhibits superior performance on app identification and activity recognition with an accuracy of more than 90%, compared to the LSTM and

**Figure 24: Accuracy of app identification and activity recognition over different VR headsets.****Figure 25: Accuracy of app identification and activity recognition across different numbers of bands.**

transformer models with an accuracy of less than 90%. The transformer model performs better than LSTM due to the self-attention mechanism that can capture emanations over frequency and time. The superior performance of ResNet is mainly due to its revolutionary architecture, as it leverages residual learning to address the vanishing gradient problem of DNN. As such, this design is particularly effective for extracting and classifying emanation features amidst diverse indoor environments.

7.3.3 Impact of the number of frequency bands. Since the emanations spread across a wide frequency band, it is important to extract all the emanation spikes for accurate app identification and activity recognition. So, we need to explore the impact of the number of bands on app identification and activity recognition.

Method. To do so, we mainly explore the frequency band between 585MHz and 626MHz with a bandwidth of 10MHz. As such, we have five frequency bands as the input of the fine-tuned ResNet for app identification and activity recognition. We report the accuracy of app identification and activity recognition when we use different numbers of frequency bands.

Result. Fig. 25 shows the accuracy of the VR app identification and activity recognition over different numbers of frequency bands. As we can see, when there is only one frequency band, the accuracy of app identification is around 0.8. However, as the number of frequency bands increases, the accuracy of app identification increases to around 0.98. Since the app identification accuracy approaches one, we cannot further increase the accuracy when we use more than two frequency bands. This is because different apps can be accurately characterized by the emanation spikes that are periodically spread across a wide frequency band. This also indicates why we focus on frequency bands below 1 GHz. However, an app activity recognition accuracy is around 0.99, even with the emanations from one frequency band. This is because we leverage the over-time frequency characterization of the emanations to differentiate app activities.

7.3.4 Impact of emanation signal duration. Emanation signals are sniffed with the software-defined radios over a period of time duration that can affect the frequency-domain emanation signal characterization.

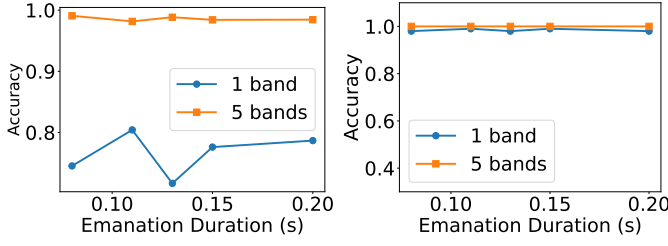


Figure 26: Accuracy of app identification across different emanation durations.

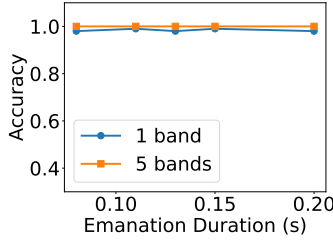


Figure 27: Accuracy of app activity recognition over emanation durations.

Method. To measure the impact of emanation signal duration on app identification and activity recognition, we derive the frequency-domain emanation representation over different emanation signal durations. The emanation signals are eavesdropped with USRP N210 at a sampling rate of 25 MHz for fine-grained spectrum characterization.

Result. Fig. 26 shows the accuracy of identifying the VR apps across different emanation durations. As we can see, the accuracy is around 0.98 across different emanation durations when we rely on emanations across five frequency bands. However, the accuracy is reduced to 0.76 across different emanation durations when we only rely on emanations over one frequency band. The accuracy does not vary significantly over the different durations. This is because the emanations are mainly characterized across frequency bands. The time-domain characteristics affected by the noise can be mitigated by our subtraction approach and the capability of the ResNet model. Moreover, the emanations are amplitude-modulated clock signals, whereby the emanation spikes can spread across a wide frequency band. As such, app identification across more frequency bands could provide better performance.

Fig. 27 shows the accuracy of app activity recognition across different emanation durations. As we can see, the accuracy of app activity recognition is around 0.99 when we rely on emanations spreading across one or five frequency bands for app activity recognition. This is because we use over-time frequency characterization of the emanations for app activity recognition, which aligns with what we have discussed about the impact of the number of frequency bands. Moreover, the accuracy does not change significantly over different durations due to the high sampling rate of 25 MHz, which can capture the detailed emanation characteristics.

7.3.5 Impact of distance. The emanations can be attenuated over the air, which can affect the emanation reception strength at the eavesdropper. Intuitively, the longer the traversing distance, the weaker the reception of the emanation at the eavesdropper.

Method. To evaluate the impact of the distance between the eavesdropper and the VR user’s headset, we vary this distance to see the changes in the emanation spikes. We fixed the directional antenna’s orientation at 90 degrees to the VR headset as shown in Fig. ?? . As we can see in Fig. 12, the USNR of the emanation spikes decreases over increasing distances. However, in this experimental evaluation, we plan to focus on the USNR variations of emanation spikes and

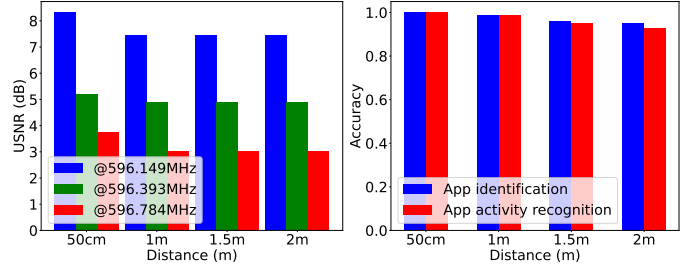


Figure 28: USNR of the emanation spikes over different distances.

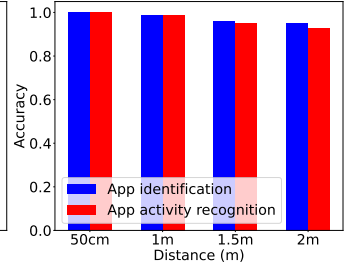


Figure 29: Accuracy of app identification and activity recognition over distances.

the accuracy of app identification and activity recognition over longer distances.

Result. Fig. 28 shows the USNR of the emanation spikes over the distances at frequencies of 596.149 MHz, 596.393 MHz, and 596.784 MHz. As we can see, the USNR decreases as the distance increases due to the over-the-air emanation attenuation. When the distance is over 1m, the USNR of the emanation spikes at the frequency of 596.149 MHz is around 7.5 dB. This is because the multipath effect becomes dominant over a longer distance, which can strengthen the received emanation signals due to the constructive signal addition. Moreover, our subtraction approach can eliminate ambient wireless interference to enhance emanation detection. We also notice that the strength of the emanation spike decreases as the frequency increases, which complies with the characteristics of the spreading emanation signals over a wide frequency band, as we have discussed in the background section. To further improve the emanation strength, we can add a power amplifier at the eavesdropper. We also notice that different emanation spikes exhibit different strengths due to the spreading spectrum property of the emanations. Fig. 29 shows the accuracy of app identification and activity recognition over different distances. As we can see, the accuracy of app identification and activity recognition slightly decreases as the distance increases. This is because the longer distance results in weaker emanation reception, which can be easily distorted by the ambient interference.

7.3.6 Impact of eavesdropper-headset orientation. Since the eavesdropper uses the directional antenna to maximize the emanation signal reception, the orientation of the directional antenna to the VR headset could affect the emanation reception strength. Therefore, we evaluate the impact of orientation on the USNR of the emanation spikes.

Method. The eavesdropper’s directional antenna can face the VR headset at different orientations. We mainly consider eight different orientations from 45° to 360° with 45° separation. When the directional antenna faces the front end of the VR headset, it is 90 degrees. When the directional antenna faces the back end of the VR headset, it is 270 degrees. In this case, the emanations should penetrate the victim’s head or be reflected off the objects in the environment for proper emanation sniffing. Note that the distance between the headset and the directional antenna is 1 meter across all the orientations. Then, we sniff the emanations to derive the

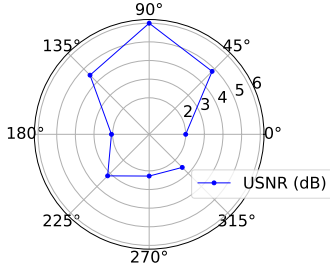


Figure 30: USNR of the emanation spikes when the eavesdropper’s directional antenna faces the VR headset at different orientations.

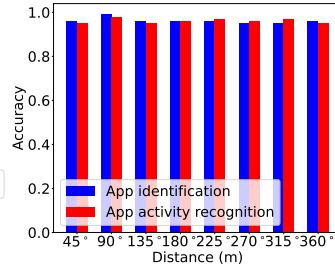


Figure 31: Accuracy of the app identification and activity recognition over different eavesdropper-headset orientations.

USNR of the emanation spikes. As we can see, when the directional antenna faces the VR headset’s front end, it is 90 degrees. We should receive the emanation signals with maximum strength. We focus on the emanation spikes within the frequency band between 590 MHz and 595 MHz.

Result. Fig. 30 shows the USNR of the emanation spikes when the eavesdropper’s directional antenna faces the VR headset at different orientations. As we can see, when the directional antenna faces the left and right sides of the VR headset, the emanation strength becomes the least (i.e., less than 1 dB) due to the poor directional antenna’s orientation. When the directional antenna’s orientation is between 45° and 135°, the USNR becomes larger. At the orientation of 90°, the USNR is maximized. However, when the directional antenna’s orientation is between 180° and 360°, the emanation strength becomes weaker due to the head blockage. We find that the emanation strength becomes larger when the directional antenna is facing the front end of the VR headset. This is because the VR headset is worn on the victim’s head. As such, when the directional antenna is facing the back end of the VR headset, the emanations are attenuated by the victim’s head. Therefore, the attacker can use the directional antenna facing the VR headset’s front end to maximize the emanation reception performance. Fig. 31 shows the accuracy of the app identification and activity recognition when the eavesdropper’s directional antenna faces different orientations to the VR headset. As we can see, the accuracy of app identification and activity recognition is around 0.96 across the orientations. This is because VReaves leverages the machine learning models to characterize the signal pattern of the emanations that are not affected by the orientations. Moreover, even though the orientations could affect the USNR of the emanations, VReaves’s system performance is not affected as long as the emanation spikes are characterized.

8 Discussion

8.1 Countermeasures

EM prevention and jamming. To defend against emanation-based privacy attacks on the VR headset, the straightforward idea is to prevent the emanation leakage from the VR headset through

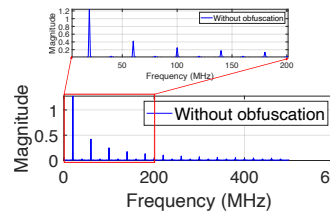


Figure 32: Frequency-domain emanations without EM obfuscation.

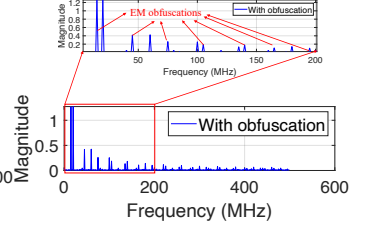


Figure 33: Frequency-domain emanations with EM obfuscation.

shielding or jamming the eavesdroppers. However, this is fundamentally difficult, as the emanations are automatically and unintentionally emitted from the VR headset, consisting of different kinds of IoT devices such as cameras and microphones. Shielding cannot fully prevent the emanation leakage. For example, we cannot fully shield the camera lens. Moreover, the jamming across the spectral spreading emanations could interfere with the in-progress legitimate wireless communication without knowing the eavesdropper’s location.

EM obfuscation. Another method to defend against this emanation-based attack on the VR headset is to obfuscate the emanations when they are emitted from the VR headset. As we have illustrated before, the emanations are amplitude-modulated clock signals that are affected by the computational activities on the VR headset. So, if we can obfuscate the computational activities on the CPU of the VR headset by running a daemon program, we can disable the emanation-based VR app identification and activity recognition. As shown in Fig. 32, we simulate the squared waves to represent the emanations in the frequency domain that are emitted from running VR apps without obfuscation. Fig. 33 shows the frequency-domain emanations when the obfuscation is added through the simulated squared waves that can represent the emanations from running the daemon in the background. As we can see, the spectrum becomes different, and the simulated obfuscations are hard to suppress without prior knowledge of them, especially when the obfuscations are random. However, this can add extra overhead to the VR headset due to the daemon running in the background.

Inference obfuscation. Since our goal is to build the relationship between the emanations from the VR headset and the VR app identity and activity with machine learning models, we can obfuscate the machine learning model to disable this attack through adversarial learning, which highly relies on intelligent EM obfuscation in the above to generate adversarial examples in reality.

8.2 Limitations and Future Work

Beyond app information inference. In our study, we mainly focus on inferring the VR app characteristics. The emanations emitted from the VR headset can also reveal other important information. For example, we can use the emanations to infer the video contents and even reconstruct the video scenes in the VR headset, which requires us to build a relationship between the emanations and video contents with advanced machine learning models. This is because the video contents are related to the computational activities on the

CPU of the VR headset, which can be revealed through emanations. Moreover, we focus on VR apps that are popular on the Meta Quest 3 and HTC VIVE XR Elite platforms. We believe that more types of apps and activities can still be accurately identified and recognized with well-trained machine learning models proposed in our work.

Multi-user VR. In our work, we mainly exploit the emanations to reveal the VR app characteristics from a single VR user for the targeted attack. Specifically, the attacker can use a directional antenna to target a specific VR user for app identification and activity recognition. When there are multiple VR headsets running different VR apps in a multi-user VR scenario, it is feasible to direct the directional antenna for emanation sniffing. However, it is difficult to differentiate the emanations from multiple VR headsets simultaneously. This is because the emanations from them spread across a wide frequency band and interleave with each other. The possible solution is to leverage the hardware imperfections of the VR headsets to differentiate the emanations from multiple headsets, which requires the VR headset to be characterized experimentally.

Long-range eavesdropping. In our work, we boost the emanation strength through the directional antenna and subtraction approach. To further boost the emanation strength, we can leverage electromagnetic interference (EMI) indicated in DeHiREC [67]. The basic idea of using EMI to boost the emanation strength is that the actively transmitted EMI could resonate with the hardware components in the VR headset to generate stronger emanations. To do so, we need to experimentally find out the EMI frequency that can efficiently excite the VR headset and deploy an extra transmitter to inject the EMI. Moreover, in our current work, we focus on the emanations from the VR headset with a directional antenna. To further enhance the system performance, we can also eavesdrop on the emanations from the hand controller as a complement to our existing settings.

9 Related Work

Virtual reality security. Recently, virtual reality platforms have been widely studied to explore their security issues. However, the existing side-channel attacks on VR devices mainly focus on inferring the keystrokes [1, 24, 29–31, 34, 47, 58, 60, 61, 66, 66], breathing and heartbeat patterns [65], VR user’s location [10], facial muscle vibration-based authentication [65], motion-based VR user identification [35] and speech extraction [5], and network traffic-based keylogging attack [51] when the VR user is wearing the head-mounted display. These side channels include the acoustic emanations, the VR user’s behavioral information (e.g., head movements), wireless channel state information, etc. For example, VR-Spy [1] leverages channel state information (CSI) from the wireless signals to infer keystrokes in the VR headset. Recently, Heimdall [30] exploits the acoustic emanations from the VR controller to infer the keystrokes. INTRUDE [36] proposes to use the head movement information to infer the video types in VR headsets. A remote keylogging attack [51] is proposed to infer user-typed secrets using the network traffic side channel in multi-user VR applications. Papers in [47] and [36] explore the head motions to predict the user’s keylogging information and VR videos, respectively.

Unlike these works, our work mainly leverages the automatic and unintentional electromagnetic emanations from the VR headset to infer the VR app identities and activities, which have never been explored in prior works. Moreover, it is not clear that the previously explored side-channel information (e.g., head motions) could be related to the VR app identities and app activities. However, the emanations can reveal the computational activities of the VR headset, which can be used for app identification and activity recognition. The prior works mainly focus on accurately extracting the frequency of the emanation spikes from the individual emanation source (e.g., camera, microphone) without considering the power spectral density and the over-time frequency properties of the emanations. Our work leverages the machine learning model to characterize the emanations from multiple emanation sources in the VR headset over time and frequency properties accurately.

Electromagnetic side channel. The electromagnetic side channel (i.e., emanation) has been studied for privacy attacks recently. Basically, the adversaries mainly eavesdrop on the side-channel information of the emanations emitted from the electronic device to infer private information such as keystrokes [20, 57, 59], secret keys [9, 12–14, 16], video, image, or audio contents [7, 8, 17, 18, 22, 26, 27, 38, 56, 63, 67]. However, these prior works have not exploited the characterized emanations for VR app identification and activity recognition. Moreover, these emanation-based privacy attacks do not fundamentally eliminate the interference from the ambient wireless communication signals. For example, EM Eye [28] infers the video content from the monitors or cameras. EM Eye mainly characterizes the emanations in the less-crowded spectrum using the amplifier to boost the emanation strength. Moreover, the emanations emitted from the large monitors are usually strong. The emanations can also be used to detect the hidden or concealed IoT devices in the indoor environment for privacy protection [6, 23, 25, 45, 50]. For example, RFScan [52] characterizes the emanations emitted from the hidden IoT devices for detection, fingerprinting, and localization, which mainly relies on signal processing to characterize the emanations. IoTProsector [53] designs a side-channel information-based inference approach and interactive tool for IoT enthusiasts to debug IoT devices based on the emanations and other side-channel information. However, IoTProsector mainly attaches an EM probe to the IoT devices for emanation detection and characterization.

Unlike previous works, our work moves one step further to reveal the characteristics of VR apps based on the emanations. Moreover, we cannot simply employ the prior techniques for VR identification and activity recognition due to domain-related concerns, which require over-time frequency-domain emanation characterizations and efficient machine learning models.

10 Conclusion

In this paper, we thoroughly exploit the emanations emitted from the VR headset to infer the VR app identities and activities. To do so, we design VReaves, a system that can enhance emanation strength, suppress ambient wireless interference, and further characterize emanations with machine learning models for VR app identification and activity recognition. Our experimental evaluation reveals the efficiency of VReaves. We believe VReaves marks the first step in emanation-based information inference in VR.

References

- [1] Abdullah Al Arafat, Zhishan Guo, and Amro Awad. 2021. Vr-spy: A side-channel attack on virtual key-logging in vr headsets. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. IEEE, 564–572.
- [2] Apple. 2025. Apple vision pro. <https://www.apple.com/apple-vision-pro/>
- [3] Oluleke Bamodu and Xu Ming Ye. 2013. Virtual reality and virtual reality system components. *Advanced materials research* 765 (2013), 1169–1172.
- [4] Giovanni Camurati, Sebastian Poeplau, Marius Muench, Tom Hayes, and Aurélien Francillon. 2018. Screaming channels: When electromagnetic side channels meet radio transceivers. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 163–177.
- [5] Derin Cayir, Reham Mohamed, Riccardo Lazzeretti, Marco Angelini, Abbas Acar, Mauro Conti, Z Berkay Celik, and Selcuk Uluagac. 2025. Speak Up, I'm listening: Extracting speech from zero-permission VR sensors. In *NDSS*.
- [6] Anadi Chaman, Jiaming Wang, Jiachen Sun, Haitham Hassanieh, and Romit Roy Choudhury. 2018. Ghostbuster: Detecting the presence of hidden eavesdroppers. In *Proceedings of the 24th annual international conference on mobile computing and networking*. 337–351.
- [7] Huiling Chen, Wenqiang Jin, Yupeng Hu, Zhenyu Ning, Kenli Li, Zheng Qin, Mingxing Duan, Yong Xie, Daibo Liu, and Ming Li. 2024. Eavesdropping on Black-box Mobile Devices via Audio Amplifier's EMR. In *Proceedings of the 2018 Annual International Conference on Network and Distributed System Security (NDSS)*.
- [8] Jieun Choi, Hae-Yong Yang, and Dong-Ho Cho. 2020. Tempest comeback: A realistic audio eavesdropping threat on mixed-signal socs. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 1085–1101.
- [9] Jesse De Meulemeester, Antoon Purnal, Lennert Wouters, Arthur Beckers, and Ingrid Verbauwhede. 2023. {SpectrEM}: Exploiting Electromagnetic Emanations During Transient Execution. In *32nd USENIX Security Symposium (USENIX Security 23)*. 6293–sun2025revealing6310.
- [10] Habiba Farrukh, Reham Mohamed, Aniket Nare, Antonio Bianchi, and Z Berkay Celik. 2023. {LocIn}: Inferring semantic location from spatial maps in mixed reality. In *32nd USENIX Security Symposium (USENIX Security 23)*. 877–894.
- [11] Gonzalo Munilla Garrido, Vivek Nair, and Dawn Song. 2023. Sok: Data privacy in virtual reality. *arXiv preprint arXiv:2301.05940* (2023).
- [12] Daniel Genkin, Noam Nissan, Roei Schuster, and Eran Tromer. 2022. Lend Me Your Ear: Passive Remote Physical Side Channels on {PCs}. In *31st USENIX Security Symposium (USENIX Security 22)*. 4437–4454.
- [13] Daniel Genkin, Lev Pachmanov, Itamar Pipman, and Eran Tromer. 2015. Stealing keys from PCs using a radio: Cheap electromagnetic attacks on windowed exponentiation. In *Cryptographic Hardware and Embedded Systems—CHES 2015: 17th International Workshop, Saint-Malo, France, September 13–16, 2015, Proceedings 17*. Springer, 207–228.
- [14] Gabriel Goller and Georg Sigl. 2015. Side channel attacks on smartphones and embedded devices using standard radio equipment. In *International Workshop on Constructive Side-Channel Analysis and Secure Design*. Springer, 255–270.
- [15] Sindhu Reddy Kalathur Gopal, Diksha Shukla, James David Wheelock, and Nitesh Saxena. 2023. Hidden reality: Caution, your hand gesture inputs in the immersive virtual world are visible to all! In *32nd USENIX security symposium (USENIX Security 23)*. 859–876.
- [16] Mordechai Guri, Matan Monitz, and Yuval Elovici. 2016. USBee: Air-gap covert-channel via electromagnetic emission from USB. In *2016 14th Annual Conference on Privacy, Security and Trust (PST)*. IEEE, 264–268.
- [17] Yuichi Hayashi, Naofumi Homma, Mamoru Miura, Takafumi Aoki, and Hideaki Sone. 2014. A threat for tablet pcs in public space: Remote visualization of screen images using em emanation. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 954–965.
- [18] Yu-ichi Hayashi, Naofumi Homma, Yohei Toriumi, Kazuhiro Takaya, and Takafumi Aoki. 2016. Remote visualization of screen images using a pseudo-antenna that blends into the mobile environment. *IEEE Transactions on Electromagnetic Compatibility* 59, 1 (2016), 24–33.
- [19] HTV. 2025. HTV vive. <https://www.vive.com/us/>
- [20] Wenqiang Jin, Srinivasan Murali, Huadi Zhu, and Ming Li. 2021. Periscope: A keystroke inference attack using human coupled electromagnetic emanations. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. 700–714.
- [21] Silvia Erika Kober and Christa Neuper. 2013. Personality and presence in virtual reality: Does their relationship depend on the used presence measure? *International Journal of Human-Computer Interaction* 29, 1 (2013), 13–25.
- [22] 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.
- [23] Yue Li, Zhenxiong Yan, Wenqiang Jin, Zhenyu Ning, Daibo Liu, Zheng Qin, Yu Liu, Huadi Zhu, and Ming Li. 2024. GPSBuster: Busting out Hidden GPS Trackers via MSoC Electromagnetic Radiations. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*. 3302–3316.
- [24] Zhen Ling, Zupei Li, Chen Chen, Junzhou Luo, Wei Yu, and Xinwen Fu. 2019. I know what you enter on gear vr. In *2019 IEEE Conference on Communications and Network Security (CNS)*. IEEE, 241–249.
- [25] Ziwei Liu, Feng Lin, Chao Wang, Yijie Shen, Zhongjie Ba, Li Lu, Wenyao Xu, and Kui Ren. 2023. Camradar: Hidden camera detection leveraging amplitude-modulated sensor images embedded in electromagnetic emanations. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (2023), 1–25.
- [26] Zhuoran Liu, Niels Samwel, Leo Weissbart, Zhengyu Zhao, Dirk Lauret, Lejla Batina, and Martha Larson. 2020. Screen gleaning: A screen reading TEMPEST attack on mobile devices exploiting an electromagnetic side channel. *arXiv preprint arXiv:2011.09877* (2020).
- [27] Yan Long, Qinhong Jiang, Chen Yan, Tobias Alam, Xiaoyu Ji, Wenyuan Xu, and Kevin Fu. [n. d.]. EM Eye: Characterizing Electromagnetic Side-channel Eavesdropping on Embedded Cameras. ([n. d.]).
- [28] Yan Long, Qinhong Jiang, Chen Yan, Tobias Alam, Xiaoyu Ji, Wenyuan Xu, and Kevin Fu. 2024. EM Eye: Characterizing Electromagnetic Side-channel Eavesdropping on Embedded Cameras. *Proceedings of ACM NDSS* (2024).
- [29] Shiqing Luo, Xinyu Hu, and Zhisheng Yan. 2022. Holologger: Keystroke inference on mixed reality head mounted displays. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 445–454.
- [30] Shiqing Luo, Anh Nguyen, Hafsa Farooq, Kun Sun, and Zhisheng Yan. 2024. Eavesdropping on controller acoustic emanation for keystroke inference attack in virtual reality. In *The Network and Distributed System Security Symposium (NDSS)*.
- [31] Shiqing Luo, Anh Nguyen, Chen Song, Feng Lin, Wenyao Xu, and Zhisheng Yan. 2020. OcuLock: Exploring human visual system for authentication in virtual reality head-mounted display. In *2020 Network and Distributed System Security Symposium (NDSS)*.
- [32] Matlab. 2025. Movmedian filter for noise revmval. <https://de.mathworks.com/help/matlab/ref/movmedian.html>
- [33] Meta. 2025. Meta Quest. https://www.meta.com/quest/?srsltid=AfmBOonvAwvK-u3NufVs2_NCZdvWczg0zW84XtkdjQwwsHEi54sB
- [34] Ulkü Meteriz-Yıldiran, Necip Fazıl Yıldiran, Amro Awad, and David Mohaisen. 2022. A keylogging inference attack on air-tapping keyboards in virtual environments. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 765–774.
- [35] Vivek Nair, Wenbo Guo, Justus Mattern, Rui Wang, James F O'Brien, Louis Rosenberg, and Dawn Song. 2023. Unique identification of 50,000+ virtual reality users from head & hand motion data. *arXiv. Preprint posted online on Feb 17* (2023).
- [36] Anh Nguyen, Xiaokuan Zhang, and Zhisheng Yan. 2024. Penetration Vision through Virtual Reality Headsets: Identifying 360-degree Videos from Head Movements. *CoRR* (2024).
- [37] Tao Ni, Guohao Lan, Jia Wang, Qingchuan Zhao, and Weitao Xu. 2023. Eavesdropping mobile app activity via {Radio-Frequency} energy harvesting. In *32nd USENIX Security Symposium (USENIX Security 23)*. 3511–3528.
- [38] Tao Ni, Xiaokuan Zhang, and Qingchuan Zhao. 2023. Recovering fingerprints from in-display fingerprint sensors via electromagnetic side channel. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 253–267.
- [39] Federica Pallavicini, Alessandro Pepe, and Maria Eleonora Minissi. 2019. Gaming in virtual reality: What changes in terms of usability, emotional response and sense of presence compared to non-immersive video games? *Simulation & Gaming* 50, 2 (2019), 136–159.
- [40] Ken Pfeuffer, Matthias J Geiger, Sarah Prange, Lukas Mecke, Daniel Buschke, and Florian Alt. 2019. Behavioural biometrics in vr: Identifying people from body motion and relations in virtual reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [41] Atieh Poushneh. 2018. Augmented reality in retail: A trade-off between user's control of access to personal information and augmentation quality. *Journal of Retailing and Consumer Services* 41 (2018), 169–176.
- [42] Ettus research. 2025. LP0410 Antenna. <https://www.ettus.com/all-products/lp0410/>
- [43] Ettus research. 2025. USRP N210. <https://www.ettus.com/all-products/>
- [44] Ettus research. 2025. VERT900 Antenna. <https://www.ettus.com/all-products/vert900/>
- [45] Cheng Shen and Jun Huang. 2021. {EarFisher}: Detecting Wireless Eavesdroppers by Stimulating and Sensing Memory {EMR}. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*. 873–886.
- [46] Cheng Shen, Tian Liu, Jun Huang, and Rui Tan. 2021. When LoRa meets EMR: Electromagnetic covert channels can be super resilient. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1304–1317.
- [47] Carter Slocum, Yicheng Zhang, Nael Abu-Ghazaleh, and Jiasi Chen. 2023. Going through the motions: {AR/VR} keylogging from user head motions. In *32nd USENIX Security Symposium (USENIX Security 23)*. 159–174.
- [48] Carter Slocum, Yicheng Zhang, Erfan Shayegani, Pedram Zaree, Nael Abu-Ghazaleh, and Jiasi Chen. 2024. That Doesn't Go There: Attacks on Shared

- State in {Multi-User} Augmented Reality Applications. In *33rd USENIX Security Symposium (USENIX Security 24)*. 2761–2778.
- [49] Sony. 2025. Sony Playstation. <https://direct.playstation.com/en-us/buy-consoles/playstationvr2>
- [50] Colin Stagner, Andrew Conrad, Christopher Osterwise, Daryl G Beetner, and Steven Grant. 2011. A practical superheterodyne-receiver detector using stimulated emissions. *IEEE Transactions on Instrumentation and Measurement* 60, 4 (2011), 1461–1468.
- [51] Zihao Su, Kunlin Cai, Reuben Beeler, Lukas Dresel, Allan Garcia, Ilya Grishchenko, Yuan Tian, Christopher Kruegel, and Giovanni Vigna. 2024. Remote Keylogging Attacks in Multi-user {VR} Applications. In *33rd USENIX Security Symposium (USENIX Security 24)*. 2743–2760.
- [52] Wei Sun, Hadi Givvehchian, and Dinesh Bharadia. 2025. Revealing Hidden IoT Devices through Passive Detection, Fingerprinting, and Localization. *Proceedings on Privacy Enhancing Technologies* (2025).
- [53] Wei Sun, Yuwei Xiao, Haojian Jin, and Dinesh Bharadia. 2023. On the Feasibility of Reasoning about the Internal States of Blackbox IoT Devices Using Side-Channel Information. *arXiv preprint arXiv:2311.13761* (2023).
- [54] Sasha Targ, Diogo Almeida, and Kevin Lyman. 2016. Resnet in resnet: Generalizing residual architectures. *arXiv preprint arXiv:1603.08029* (2016).
- [55] Anastasios Theodoropoulos and Angeliki Antoniou. 2022. VR games in cultural heritage: A systematic review of the emerging fields of virtual reality and culture games. *Applied Sciences* 12, 17 (2022), 8476.
- [56] Toshihide Tosaka, Kazumasa Taira, Yukio Yamanaka, Atsuhiko Nishikata, and Mitsuo Hattori. 2006. Feasibility study for reconstruction of information from near field observations of the magnetic field of laser printer. In *2006 17th International Zurich Symposium on Electromagnetic Compatibility*. IEEE, 630–633.
- [57] Martin Vuagnoux and Sylvain Pasini. 2009. Compromising electromagnetic emanations of wired and wireless keyboards.. In *USENIX security symposium*, Vol. 8. 1–16.
- [58] Hanqiu Wang, Zihao Zhan, Haoqi Shan, Siqi Dai, Maximilian Panoff, and Shuo Wang. 2024. GAZEexploit: Remote Keystroke Inference Attack by Gaze Estimation from Avatar Views in VR/MR Devices. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*. 1731–1745.
- [59] Litao Wang and Bin Yu. 2011. Analysis and measurement on the electromagnetic compromising emanations of computer keyboards. In *2011 Seventh International Conference on Computational Intelligence and Security*. IEEE, 640–643.
- [60] Yi Wu, Cong Shi, Tianfang Zhang, Payton Walker, Jian Liu, Nitesh Saxena, and Yingying Chen. 2023. Privacy leakage via unrestricted motion-position sensors in the age of virtual reality: A study of snooping typed input on virtual keyboards. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3382–3398.
- [61] Zhuolin Yang, Zain Sarwar, Iris Hwang, Ronik Bhaskar, Ben Y Zhao, and Haitao Zheng. 2024. Can virtual reality protect users from keystroke inference attacks?. In *33rd USENIX Security Symposium (USENIX Security 24)*. 2725–2742.
- [62] Sihan Yu, Jingjing Fu, Chenxu Jiang, Chunhui Lin, Zhenkai Zhang, Long Cheng, Ming Li, Xiaonan Zhang, and Linke Guo. 2024. FreeEM: Uncovering Parallel Memory EMR Covert Communication in Volatile Environments. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*. 372–384.
- [63] Zihao Zhan, Zhenkai Zhang, Sisheng Liang, Fan Yao, and Xenofon Koutsoukos. 2022. Graphics peeping unit: Exploiting em side-channel information of gpus to eavesdrop on your neighbors. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1440–1457.
- [64] Qibo Zhang, Daibo Liu, Xinyu Zhang, Zhichao Cao, Fanzi Zeng, Hongbo Jiang, and Wenqiang Jin. 2024. Eye of Sauron: {Long-Range} Hidden Spy Camera Detection and Positioning with Inbuilt Memory {EM} Radiation. In *33rd USENIX Security Symposium (USENIX Security 24)*. 109–126.
- [65] Tianfang Zhang, Zhengkun Ye, Ahmed Tanvir Mahdad, Md Mojibur Rahman Redoy Akanda, Cong Shi, Yan Wang, Nitesh Saxena, and Yingying Chen. 2023. FaceReader: unobtrusively mining vital signs and vital sign embedded sensitive info via AR/VR motion sensors. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 446–459.
- [66] Yicheng Zhang, Carter Slocum, Jiasi Chen, and Nael Abu-Ghazaleh. 2023. It's all in your head (set): Side-channel attacks on {AR/VR} systems. In *32nd USENIX Security Symposium (USENIX Security 23)*. 3979–3996.
- [67] Ruochen Zhou, Xiaoyu Ji, Chen Yan, Yi-Chao Chen, Wenyuan Xu, and Chaohao Li. 2023. Dehirec: Detecting hidden voice recorders via adc electromagnetic radiation. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3113–3128.