

SpiralSpy: Exploring a Stealthy and Practical Covert Channel to Attack Air-gapped Computing Devices via mmWave Sensing

Zhengxiong Li*, Baicheng Chen[†], Xingyu Chen*, Huining Li[†], Chenhan Xu[†],
Feng Lin[§], Chris Xiaoxuan Lu[‡], Kui Ren[§] and Wenyao Xu[†]

* University of Colorado Denver, Denver, Colorado, USA
Email: {zhengxiong.li, xingyu.chen}@ucdenver.edu

[†] University at Buffalo, the State University of New York, Buffalo, New York, USA
Email: {baicheng.huiningl,chenhanx,wenyaoxu}@buffalo.edu

[‡] University of Edinburgh, Edinburgh, Scotland, United Kingdom
Email: {xiaoxuan.lu}@ed.ac.uk

[§] Zhejiang University, Hangzhou, Zhejiang, China
Email: {flin, kuiren}@zju.edu.cn

Abstract—Covert channels are a method of communication that is used to exfiltrate information from computing devices and break the security policy of computer systems. Any shared resource can be potentially leveraged as a covert channel, and conventional wisdom of cyber-security believes that air-gapped computing devices, disconnected from the Internet, are highly secured. Recent studies show that advanced covert channel attacks using acoustic, thermal, and electromagnetic effects can only work under a limited proximity constraint (e.g., within 2 meters). In this work, we present **SpiralSpy**, a new covert channel to attack air-gapped computing devices through millimeter-wave (mmWave) sensing technologies. **SpiralSpy** can be stealthily launched and circumvent strongly isolated computing devices from a practical distance (up to 8 meters). Specifically, we demonstrate that ordinal cooling fans can be leveraged for covert channel attacks. A malicious software inside air-gapped computing devices can saliently encode confidential data into the fan control signals, and modulated status on fan motions can be remotely decoded by a commodity mmWave sensor. **SpiralSpy** can be adopted on multiple-fan systems and enable a scalable capacity for multi-channel and high-speed information transfer. We evaluate **SpiralSpy** with 71 computing devices with cooling fans. Experimental results demonstrate that **SpiralSpy** can achieve up to 6 bps that is 6-24X faster than existing covert channels on air-gapped computing devices. We evaluate the usability and robustness of **SpiralSpy** under different real-world scenarios. Moreover, we conduct in-depth analysis and discussion on countermeasures for **SpiralSpy**-based covert channel attacks to improve computer and information security.

I. INTRODUCTION

To date, the covert channel remains a genuine threat to information flow control and isolation techniques in various

Zhengxiong Li and Baicheng Chen are co-primary authors.

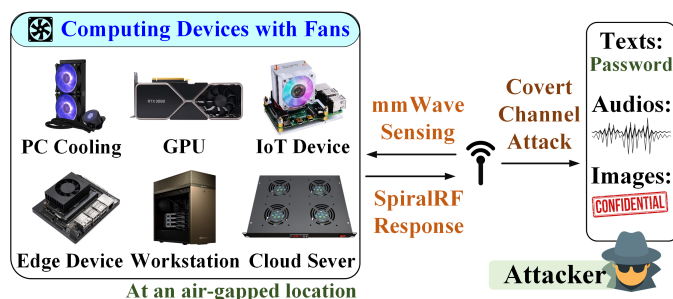


Fig. 1. SpiralSpy is a new covert channel attack that can steal confidential information from secure sites through mmWave sensing with cooling fans on various air-gapped devices.

contexts, from personal computers (PC) to cloud servers (hereafter called **computing devices**) [1]. Such channels can be maliciously exploited by attackers through data exfiltration and colluding applications [2]. Many research works have indicated that communication media, including networks [3], [4], shared cache/memory [5], [6], and I/O devices (e.g., USBee [7]), are the most vulnerable part that can be leveraged for covert channel attacks. These threats lead to the development of mitigation techniques, such as resource partition on communication media [8], which restrict dedicated resources to individual processes for the execution.

Air-gap is a fundamental mitigation to covert channel attacks by creating a strong isolation environment that excludes commodity communication media. As a practice under this concept, air-gapped computing devices are designed to stay physically isolated from insecure networks and inaccessible to unauthorized users [9]. Although the design of air-gapped computing devices represents nearly the maximum protection, *researchers and hackers have explored covert channels through nontraditional artifacts*, including acoustic signaling [10], magnetic leakage [11], [12] and optical surfaces [13]. Besides the research, the StuxNet is a famous air-gapped attack example through removable media in the real-world [14]. However, these covert channels require either proximity (e.g., within a close line-of-sight distance) or network access (e.g.,

send packets over local area network), where the mitigation can be effective through rigorous physical monitoring and local area network control. Besides, the bit rate of these covert channels is limited, which is not effective in real-world attacks. However, is it possible that a new type of covert channel exists and enables a stealthy and practical attack on air-gapped computing devices without the above limitations?

In this work, we introduce *SpiralSpy*, a fan-based new covert channel attack to air-gapped computing devices. As shown in Fig. 1, *SpiralSpy* system can steal the information from the computing devices through mmWave sensing toward cooling fans, even in an air-gapped scene in the real world. *SpiralSpy* is featured as **(1) Stealthiness**: No hardware modification, no sensor emission requirement, and no apparent abnormality; **(2) Practicality**: The attack distance can be up to 8 meters away (4-20X longer than existing non-contact covert channels) and support *through-wall* and non-line-of-sight information theft; **(3) Fast Data Rate**: It supports a high-speed data transmission that is 6-24X faster than existing non-contact covert channels and enable the practical transfer of security data files (e.g., login password, auditory key, and biometric credentials); **(4) Pervasive Media**: Opposed to primarily targeting personal computers, it can work on pervasive *computing devices* (e.g., IoT/edge devices).

The realization of *SpiralSpy* leads to the following two technical challenges: (i) How to build a new covert channel via pervasive media (i.e., fans) to retrieve confidential information from air-gapped computing devices with practical attack capabilities (e.g., long-range and through-wall)? (ii) How to design an efficient protocol to complete fast data transmission in a stealthy way?

To address these challenges, we first investigate the dependent relationship between the reflected mmWave signal and the cooling fan status within the computing devices using a portable mmWave probe. Almost every computing device today is equipped with one or multiple cooling fans. As a mmWave signal arrives at the fan, its rotational motion induces a frequency shift on the incidental wave. The resulting response (i.e., *SpiralRF* response) has a tight relationship with the fan status at the time of the incident. To extract and characterize such a frequency shift modulation, specific fan status and the resulting *SpiralRF* response are correlated using a set of 35 spectral and temporal features after wavelet transform. A random forest model is then employed to perform fan status recognition based on the extracted features. Besides, based on this relationship model, we further design and optimize a fan-based mmWave covert channel attack and information transmission protocol to solve the second challenge. Specifically, to better control fans on victim devices, an Agent sending scheme is designed and implemented for fast and steady information encoding and sending, under restrictions from low-level device drivers. Given the fact that most current computing devices are equipped with cooling fans, and there are no permission requirements to access these fans, *SpiralSpy* achieves a stealthy and practical covert channel attack towards air-gapped computing devices via mmWave sensing on the cooling fan status (e.g., fan speed). To optimize the performance of *SpiralSpy* information transfer, a multi-channel attack is designed to further improve the attack channel throughput. Subsequently, we conduct an extensive attack evaluation to

assess the performance of our system under various conditions. Eventually, we conclude the study by developing *SpiralSpy*, a practical (8m away, through-wall), low-cost, and stealthy covert channel attack that precisely steals the information from the air-gapped computing devices.

Our contribution has three-fold and summarized as follows:

- We investigate the possibility of information theft through the fan on air-gapped computing devices and elaborate a stealthy and practical covert channel without the need of modifying the computer system driver and hardware.
- We design and implement *SpiralSpy* system with an innovative information transfer (one-way communication) architecture based on the fan's *SpiralRF* response. First, we develop an Agent to plant in target devices for fan status modulation and data encoding. Then, we prototype a mmWave based scanner to interrogate and receive the *SpiralRF* response. Finally, we model and optimize the multi-channel attack with innovative command pipelining algorithms to increase communication throughput.
- We comprehensively evaluate *SpiralSpy* based on multiple affecting scenarios. First, we study the performance of our communication system, including the single and multi-channel attack accuracy. Second, we test the robustness with variances in occlusion, orientations, sensing distances, and environments. Finally, we discuss the effectiveness and study a set of countermeasures to prevent information leakage against this information threat.

II. THREAT MODEL

We consider the scenario where a target, namely Alice, implements an air-gapped protection mechanism on the computing devices (e.g., PC or IoT/edge devices) to prevent them from attacks. Observing Alice's vigilance, an attacker or law enforcement, hereafter Mallory, aims to breach the established air-gapped security and steal confidential information or collect significant evidence from these computing devices. To achieve this goal, Mallory senses the information stored in the device remotely or through-wall and reconstructs the sensitive information (e.g., login password, voice token, and credential image) without alerting Alice or the surroundings. In contrast to prior work, we envision the following constraints when Mallory practically launches an attack:

No Hardware Modification or Add-ons: Alice puts the device under a strict-controlled location so that Mallory, without permission, cannot physically access the device or add/modify the device hardware.

No Proximity: Alice is alert to traditional shoulder surfing or other line-in-sight attacks. Therefore, Mallory cannot get close to the target device, or there is physical perimeter security in place (e.g., a wall) between Mallory and the target device during the attack.

No Physical Connection: Assuming that Alice's device is physically isolated from the public Internet and well protected in the local area network, Mallory is unable to directly compromise the computing device from these network attacks.

It is worth mentioning that although it is an air-gapped computing device, there is still an inevitable need for data interaction between this device and the external storage or other

devices. Besides, infecting an air-gap computing device is well studied and can be done via many different attack vectors in an unconscious way. These unconscious ways include using targeted malicious programming tools via social engineering, planting an infected USB drive or other removable storage devices, attacking the supply chain, scheduling a malicious insider, and so on [15], [16], [17], [18]. These infections merely aim to discover and capture the information within the device while cannot send out this critical information to Mallory. *Therefore, the challenge at Mallory's hand is how to build up the novel covert channel and retrieve the stolen information from this air-gapped computing device without violating the constraints above.*

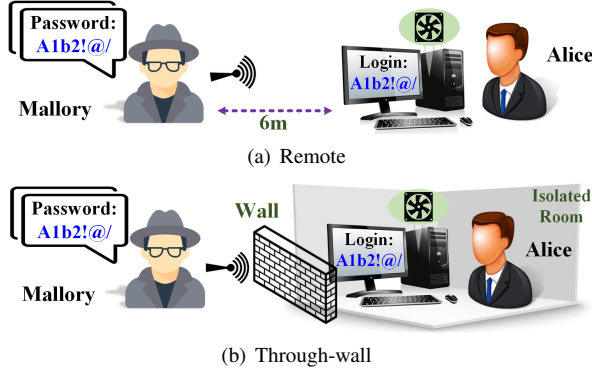


Fig. 2. Two typical attack scenarios in daily life: (a) Mallory infers the password from a remote location; (b) Mallory leverages the penetration properties of mmWave for through-wall inference.

Traditional acoustic-based, thermal-based, and EM-based covert channel attacks (e.g., [19], [10], [20]) lose effectiveness under an application scenario with the constraints mentioned above (details can be found in Section XIV). However, air-gapped computing device security in this scene will not necessarily be guaranteed if Mallory can leverage a small form-factor and low-cost mmWave probe to perform real-time surveillance of cooling fan information in an adjacent room and steal the information from the target, as shown in Fig. 2. In our work, we assume that the air-gapped computing device is equipped with cooling fan(s) and can be sensed from outside the case. Given unavoidable air vents around the cooling fans on the chassis for thermal regulation, regardless of chassis material (e.g., aluminum, tempered glass, plastic, etc.), mmWave signals can pass through, interrogate the fans, and return to the receiver.

III. BACKGROUND AND PRELIMINARIES

A. Cooling Fans in Computing Devices

The cooling fan is widely known as an essential and necessary component of the computing device. A cooling fan in computing devices consists of several fan blades that are attached to a brushless rotor. There are two mechanisms of controlling fan status (i.e., direct current-DC and pulse-width modulation-PWM), both mechanisms utilize a Hall sensor that reads real-time fan speed for accurate fan status control [21], [22]. Fans can be controlled based on these two mechanisms with a control interface in place (e.g., when manufactured) and are widely used as a thermal regulation solution in both consumer and enterprise computing devices [23]. Due to the

accessibility and ubiquity of cooling fans, it is highly plausible to utilize them for stealthy covert channel attacks on air-gapped computing devices.

B. SpiralRF Response

The key to understanding the covert channel's risk lies in accurately modeling fan status in air-gapped computing devices. To this end, we theoretically investigate the non-linear relationship between fan status and RF response. For readability, such a non-linear relationship will be dubbed *SpiralRF response* hereafter.

Fan Status in Air-Gapped Computing Devices. To remotely retrieve the fan status (or fan speed) from an air-gapped computing device, a robust physical model for sensing different cooling fans in various applications is needed. Although fan blades are often designed with a simple rectangular shape with radial symmetry, fans on computing devices often have meticulously designed fan blades to optimize various air pressure and airflow application scenarios. This essentially gives rise to the following variables in a basic fan model: (a) number of blades, (b) blade shape, and (3) rotating speed [24], [25], [26]. Concretely, given a fan with N blades, the response signal returned by the fan system s_{Σ} at a given time t with tangential velocity Ω is modelled as [27]:

$$s_{\Sigma}(t) = L * \exp\{-j \frac{4\pi}{\lambda} [R_0 + vt + y_0 \sin(\beta)]\} * \sum_{k=0}^{N-1} \sin(c\{\frac{L}{2} \frac{4\pi}{\lambda} \cos(\beta) \cos(\Omega * t + \phi_0 + 2k\pi/N)\}) * \exp\{-j \Phi_k(t)\}, \quad (1)$$

where ϕ_t is the fan blade angle at radar view at time t , Ω is the blades' instantaneous tangential velocity, β is the elevation angles of the center of the fan blades system at radar view, R_0 is the initial distance between the radar and the fan blade system, v is the fan system's radial velocity in radar direction (i.e., blade speed), L is the entire length of the blade, the blade shape can later be integrated over the length L to comply with various design, λ is the frequency of interrogation wave, P is a point on the fan blade tip, (x_0, y_0, z_0) coordinate represents the initial three-dimensional position of P at radar view, $\Phi_k(t)$ is the phase function of reflected signal at time t for fan blade at relative index k modeling as:

$$\Phi_k(t) = \frac{L}{2} \frac{4\pi}{\lambda} \cos(\beta) \cos(\Omega t + \phi_0 + 2k\pi/N), \quad (2)$$

where ϕ_0 is the initial rotation angle at time 0. At the crux, the frequency and phase of the reflected signal in the above formulation are dependent on every single fan blade's

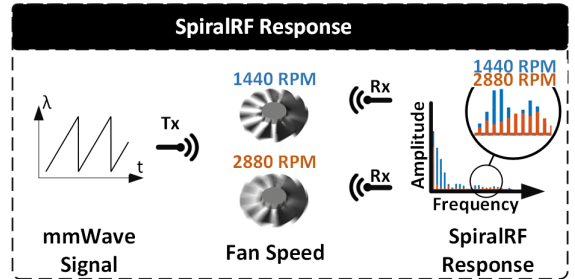


Fig. 3. *SpiralRF* responses are unique at different fan setups (e.g., 1440 rpm vs 2880 rpm) when interrogated by the same RF signal.

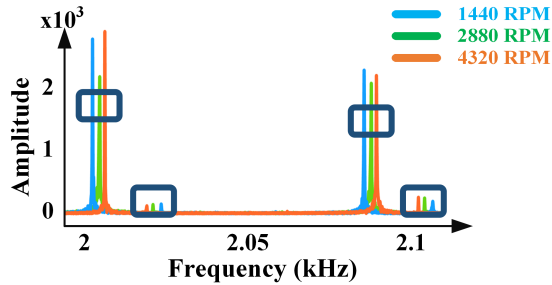


Fig. 4. The spectrum of SpiralRF responses after the signal demodulation are distinct from different fan statuses.

tangential velocity, rather than radial velocity, *which provides more fan status detail (i.e., blade shape and rotating speed) and improves sensitivity upon classic micro-Doppler effect model that can only provide general radial velocity* [26].

Although the single-frequency model is sufficient to identify the fan status in lab conditions, real-world scenarios often contain non-negligible noise from the background object’s reflection and unintended signal superposition from communication devices. To counter the impact of noise and leverage the mmWave’s wide bandwidth in radar sensing, we propose a model using frequency modulated continuous wave (FMCW) mmWave to fine-grained distinguish the fan’s reflection signal containing the fan status information, namely the **SpiralRF response**. Compared to higher frequency RF waves, the 24 GHz RF signals can penetrate the wall with better performance. On the other hand, it has a larger bandwidth (e.g., 450 MHz) to capture the subtle motions more accurately in the ambient (i.e., better than MHz RF probes with a kHz bandwidth). Thus, we utilize a 24 GHz FMCW mmWave radar with a 450 MHz bandwidth for fan speed sensing. The full reflected signal as a function of the fan’s status is given in Equation (3):

$$\begin{aligned}
 F(\Omega, t) = & \sum_{\lambda=b}^B L * \exp\{-j\frac{4\pi}{\lambda}[R_0 + vt + z_0 \sin(\beta)]\} \\
 & * \sum_{k=0}^{N-1} \sin(c\{\frac{L}{2} \frac{4\pi}{\lambda} \cos(\beta) \cos(\Omega * t + \phi_0 + 2k\pi/N)\}) \\
 & * \exp\{-j\Phi_k(t)\},
 \end{aligned} \tag{3}$$

where $F(\Omega, t)$ is the fan status dependent reflection signal from the fan, b is the initial frequency of the wireless interrogation, and B is the bandwidth. Fig. 3 demonstrates the intuition behind this signal reflection model, where we can clearly observe different fan statuses leading to different corresponding SpiralRF responses.

C. Feasibility Study

In practice, fans are often installed on the top/front/ back of the computer chassis, where large openings are made to allow sufficient airflow for cooling. Such design allows mmWave to bypass the chassis through openings and generate SpiralRF response back to the mmWave probe. Next, we examine the feasibility of wireless fan status sensing using SpiralRF response via mmWave interrogation. To this end, we sense different speeds on a PWM fan installed on the chassis of a desktop computer. A NVidia GTX Titan XP Founder Edition

GPU cooling fan (60 mm diameter with max fan speed 4800 rpm) is utilized to simulate a real-world application scenario. The fan status is then manipulated by us through the motherboard’s Unified Extensible Firmware Interface (UEFI). Given the duty cycle value (ranging from 0 to 100), the chassis fan can be manually controlled to maintain at the corresponding speed. By using a 24 GHz mmWave probe with a bandwidth of 450 MHz [28], the fan is scanned at different statuses under the commands. As shown in Fig. 4, the frequency shifted responses from a fan rotating at 1440 (30% duty cycle), 2880 (60% duty cycle), 4320 RPM (90% duty cycle) are significantly different. This feasibility test demonstrates the effectiveness of SpiralRF response in fan status sensing, which can be later utilized for constructing a fan-based covert channel attack.

IV. SPIRALSPY ARCHITECTURE

We are now in a position to introduce the main architecture of our designed system - **SpiralSpy**. At a high level, we designed an **Agent-Sniper** two-module architecture for launching through-wall covert channel attack. The Agent is responsible for continuously encoding the information into the transmittable bit data flow (0 or 1) and sending bit data through cooling fan status. On the other side, the Sniper remotely reads the cooling fans’ status and decodes the bit data after segmenting the header/ data /footer section in a packet, and then returns the binary data into corresponding information. The information attack process with the two-module architecture is shown in Fig. 5.

A. SpiralSpy Planted Agent

To steal information (uni-directional communication) from air-gapped computers, a spy malware Agent is pre-installed on computing devices to transfer the data stealthily. It is worth to mention that the Agent does not change computation performance, human perceptible noise, or cause anti-virus issues during the stealthy operation. Using a flexible baseline in fan speed parameter setting to match the operating system (OS) fan-controlling curve, **SpiralSpy** can seamlessly transmit data while effectively cooling the computing system. In addition, changes of cooling fan speed will not generate noticeable noise in human perception. Even high pitch noise from high-speed fans naturally correlates to short-term computationally intensive tasks in daily operation. Currently, hardware manufacturers [29], [30], [31] have integrated the thermal regulation process with fan speed control (FSC) modules in their drivers, and the applications can adjust the fan speed for the sake of performance/noise trade-off. Note that drivers are required to run every-time when the system starts, and it would be impractical to request user input password for drivers to run, thus. Therefore, accessibility privilege in drivers is usually granted once-for-all in the installation process. Particularly, the FSC module, being a part of the drivers, inherits such privilege and does not request additional permission when commanded to access fan speed using driver interface. This results in security vulnerabilities to **SpiralSpy** attack using the FSC module in drivers without triggering anti-virus defense.

In detail, the Agent utilizes fans on computing devices (e.g., chassis fans, GPU fans, and CPU fans) to send out the information by changing fan status to represent different information bits (e.g., 50% fan speed for 0 bit, and 51% fan




Fig. 5. The flow chart for SpiralSpy consists of an Agent-Sniper two-module architecture that allows information theft from an air-gapped computing device via mmWave sensing.

speed for 1 bit). The Agent controls fans with either DC or PWM voltage regulation, which can accelerate or decelerate the fan to a specific target level (e.g., 0% duty-cycle to 100% duty-cycle with intervals of 1%). To better manipulate the target fan status, a customized fan status control algorithm is implemented on the fan control interface, namely the Fan Speed Modulator.

B. SpiralSpy Remote Sniper

SpiralSpy Sniper comprises careful designs of hardware and software components.

Sniper Hardware: To capture each fan’s status from outside of the air-gap (e.g., few meters away from the air-gapped computer and behind the wall), the Sniper employs an FMCW mmWave radar (namely the mmWave-Spiral Sensor) that continuously interrogates the fans and then captures the corresponding SpiralRF response. As the fan status changes, its behavior under mmWave radar’s sight can be discerned based on the SpiralRF response.

Sniper Software: After receiving the SpiralRF response, the signal is processed into spatial-temporal features to efficiently describe the response signal and then be fed into a machine learning classification algorithm that quickly, accurately, and robustly recognizes the fans’ status. After that, a translation from fan status to binary data (i.e., Data Decoder) is designed and implemented to reconstruct the information according to the information transmission protocol.

V. AGENT SENDING SCHEME

In the planted Agent, Information Encoder and Fan Speed Modulator are designed to translate bits into fan speed sequences and execute the fan speed sequences accurately and timely respectively. Before the Agent transfers the data by manipulating the fan status, there are *two technical challenges* needed to be handled: (1) *Fan Hardware*: There are some limitations for fan control within the fan and its peripheral, such as around 10% inevitable errors in fan status mechanical control, 0.3s response time before fan status changing, and 0.2s delay for system notification after changing [32]. Such limitations in fan control make current one-way communication techniques, such as FSK [33], ASK [34], OFDM [35], inefficient or inadequate in this Spiral-mmWave fan covert channel. (2) *Computing Device System*: In real practice, besides the SpiralSpy manipulation, the computing device system can also manipulate the fan status simultaneously for routine thermal regulation. To ensure the attack stealthiness, SpiralSpy does not block this device system control on the fan. Instead,

SpiralSpy reads system’s target fan speed periodically and flexibly adjusts fan speed accordingly. Thereby, the challenge is to meticulously distinguish SpiralSpy manipulation of the fan from the device’s manipulation due to OS regulation.

A. Information Encoder

To address the above challenges on fan control, Information Encoder is designed to encode binary data in the victim device into a fan status sequence. In order to accurately represent binary data in fan status, a bespoke translation protocol for the Spiral-mmWave channel is designed. The Information Encoder utilizes the following three rules for packet design: (1) A SpiralSpy packet is split into header/footer, and data sections to signal the start/end and data content of the packet to be transmitted, respectively. (2) The header/footer sections aim to distinguish SpiralSpy’s fan status from the operating system’s routine fan speed adjustment, which will be abrupt for a short period. (3) The data section aims to smoothly and quietly leak information over a period of time that ensures transmission accuracy and stealthy for the covert channel.

Header and Footer Sections: The header section is responsible for marking the beginning of a packet for the receiver to decode. As shown in Fig. 6, as the time changes, the fan speed moves to a much lower value with a steep line on the speed curve and then quickly returns to the baseline fan speed (i.e., 50%) to get ready for data transmission. Such abrupt change in fan speed indicates that the Agent starts to transmit packet via this fan instead of a device system-controlled fan speed. Because of fan speed control instability, the fan speed can easily overshoot the baseline, resulting in an error in the first few transmission bits. Thus, a period is padded to increase stability and mark the first transmission bit for distinguishing between the header section and data section.




Fig. 6. SpiralSpy information transmission protocol for single fan transmission, a full packet consisting of header/data/footer section is shown.

It is also crucial to mark the ending of a packet to ensure no additional bits are added, which may cause unreadable packet content. Thus, the footer section is implemented to distinguish itself from the following OS-controlled fan speeds. Similar to the header section, the footer moves to a fan speed with a steep line on the speed curve and then quickly returns to the baseline fan speed, with the only difference being the direction in the fan speed axis, as observed in Fig. 6. However, when the footer section starts accelerating with the last bit being a one, it is likely to be ignored due to its position as part of the footer section in the fan speed plot. Thereby, a period is also padded before the footer starts accelerating to guarantee the integrity of the data section.

Data Section: The data section is designed to successfully transmit the binary data through fan speed sequence as fast as possible with integrity. Thus, to maximize the performance, a protocol with 2 bps bit rate is designed, resulting from a 2 Hz fan speed control rate and 1 bit per sample fan-speed bandwidth. To represent successive ones and zeros without constant acceleration to exceed fan speed limit, fan speed edge detection is adopted. The resulting data section protocol can be summarized as follows: **1)** We define a baseline speed f_0 and a high speed f_1 representing zero and one bits, respectively. **2)** After the padding time past header section, t_0 , the first bit b_0 with timestamp t_1 is sampled 0.5s after t_0 . **3)** For the rest of $n - 1$ bits in the data section, if the bit b_x is equal to b_{x-1} , the fan speed remains still for the next sampling duration (i.e., 0.5s). **4)** At timestamp t_n , all n bits are transmitted, the fan speed returns to baseline speed f_0 to get ready for the footer section. To this end, the protocol packet that contains a header marking the beginning of the data section and a footer marking the data section is fully functional for *SpiralSpy* covert channel.

B. Fan Speed Modulator

The responsibility for the Fan Speed Modulator is to apply the commands received from Information Encoder on the fan. However, there are many limitations on fan execution from the driver and the hardware layers, which prevents the fan speed from changing swiftly (the rate at the seconds level). Usually, an operating system’s thermal regulation module updates the fan speed on a default fixed time interval (e.g., 5s, 10s, 30s, 1 min). The native fan controlling driver software then adjusts fan speed to a new target speed taking approximately ten to twenty seconds. Such acceleration/deceleration duration will inevitably make the covert channel impractical (40s for one bit). Thus, the main challenge is to increase the data transmission rate (via rapidly changing fan speeds) with stability (into the sub-second level). Using a fan speed modulator, we increase the communication rate by one order of magnitude without modifying existing hardware and software infrastructure (e.g., fans inside computer chassis, operating system, device drivers, etc.).

Header and Footer Sections: A common intuition for speed control, such as vehicle cruise speed control, is to continually read the fan speed and adjust based on the speed delta. Given the large fan speed change interval in these two sections, it is possible to utilize such cruise speed control methods for execution. To guarantee the header and footer threshold is crossed with variations in acceleration timing, an iterative

TABLE I. AVERAGE ACCELERATION AND DECELERATION TIME DELTAS FROM 100 TRAILS

Section	Acceleration	Deceleration
Header	6.205 s	14.708 s
Footer	6.883 s	6.266 s
Data	0.326 s	0.302 s

speed change technique is designed as follows: **1)** Depending on the target fan speed, the Fan Speed Modulator accelerates or decelerates the fan for a duration that makes the fan speed very close to the target as shown in Table I, if it has not already crossed the threshold. **2)** Perform a fan speed read and determine whether more time is needed to get to the designated fan speed. If the threshold is crossed, return to the baseline fan speed by performing the previous step. If not, perform 1 s of acceleration/deceleration, and repeat.

Data Section: Different from the header and footer sections, the acceleration and deceleration duration in the data section are noticeably shorter and have higher error tolerance for an additional 0.2 s window. To guarantee that the data section’s acceleration and deceleration reach and does not overshoot target fan speed, a novel technique is developed. Considering fan speed delta being only 1%, the difference between actual fan speed and target fan speed is relatively small after the average acceleration/deceleration duration. Thus, a precisely timed acceleration/deceleration will bring the fan to target fan speed without the risk of overshooting. To this end, the Fan Speed Modulator achieves rapid transmission of bit data with integrity.

VI. SNIPER READING SCHEME

In order to read the fan speed remotely, we design the *Sniper* reading process as follows. First, the selected mmWave probe has a high frequency and a wider bandwidth (i.e., 24Ghz with a 450 MHz bandwidth) that is capable of capturing more details of an object’s motion. Second, we model the phenomenon of *SpiralRF* and build up a list (35) of physical features related to rotatory speeds. Compared to the spectrum, the use of our feature list adds robustness to the model while reducing computation cost. Third, the random forest can model the nonlinearity of *SpiralRF* and excel the accuracy performance through a data-driven approach.

A. Signal Preprocessing

To understand the fan speed, we first adopt empirical wavelet transform (EWT) [37] to separate *SpiralRF* response into a set of wavelets as shown in Equation (4).

TABLE II. SPIRALSPY FAN SPEED CLASSIFICATION FEATURES

Domains	Features
Spectral	Crest Factor, Flatness, Skewness, Kurtosis, RMS Amplitude, Lowest Value, 25th percentile, 50th percentile, 75th percentile, Highest Value, MFCC-15 [36]
Temporal	Lowest Value, 25th percentile, 50th percentile, 75th percentile, Highest Value, Mean Value, Standard Deviation, Kurtosis, Skewness, Flatness

$$f(t) = (W_f^\varepsilon(\hat{0}, \omega)(\hat{\phi}_1(\omega) + \sum_{n=1}^N (W_f^\varepsilon(\hat{n}, \omega)(\hat{\psi}_n(\omega))))))^\vee, \quad (4)$$

where $f(t)$ is the reconstruction of all wavelets that infinitely approaches the raw SpiralRF response signal, t is time, ω is the wavelength, which is inversely proportional to response signal wave frequency, $\phi_n(\omega)$ is the empirical scaling function, $\psi_n(\omega)$ is the empirical wavelets, N is the number of wavelets in total, $W_f^\varepsilon(0, t)$ is the approximation coefficients, and $W_f^\varepsilon(n, t)$ is the detail coefficients given by the inner products of the empirical wavelets. \wedge is the Fourier transform representation, and \vee is the inverse Fourier transform operation. To this end, the SpiralRF response signal is transformed into a series of empirical wavelets for feature extraction.

B. SpiralRF Feature Extraction

After getting the pre-processed signal, we adopt two categories of mmWave features on spectral and temporal domains to effectively describe the characteristics. As the fan blades rotate, the SpiralRF response from a fixed point at the circular perimeter can be seen as a binary value marking the presence of a fan blade at a specific timestamp. However, such representation is difficult to extract from response spectrum directly given a large amount of unrelated information (e.g., noise interferences, vibration, occlusions) in reflect RF signal. Although data-driven models, such as deep learning, could correlate such noise with fan speed and cause catastrophic results in real-world testing with deceitfully good training results. Thus, specialized physical features related to rotary speed are necessary for accurate inference. In addition, the spectrum information will be processed into 64x64 spectral data, which is computationally intensive, and limits the deployability in mobile devices in covert channel attacks. Thus, the use of our feature list adds accuracy/robustness to the model across various attack scenarios while reducing computation cost, which can facilitate real-time remote attack in mobile mmWave devices with low computation resources. Thus, we employ the resulting feature vector with a total of 35 dimensions is presented in Table II.

C. Fan Speed Identification

To accurately decode the fan speed from the feature vector, we employ a random forest classification module to predict the fan speed (i.e., duty cycle). Compared to traditional machine learning techniques such as linear regression [38] and support vector machine (SVM [39]), random forest [40] is intrinsically designed to perform multi-class classification (e.g., fan speed value from 25-100, a total of 76 classes) with better non-linear dependencies. Compared to deep learning techniques that are heavily data-driven, the random forest is an ensemble algorithm that utilizes many smaller classifiers in multiple layers, making it extremely resistant to overfitting. Thus, we utilize a random forest model that takes the feature array as input and fan speed as output. Specifically, the model contains 100 decision trees for efficiency considerations, and the out-of-bag classification error is expected to converge to zero before all 100 trees are used.

D. Fan Speed-Binary Data Translation

To achieve real-time communication, real-time fan speed decoding is performed by the Data Decoder for header section detection. Thus, a 30-second window of fan speed at a 10 Hz sampling rate is kept in storage to check if the header section pattern occurs. *To prevent false alarm*, the baseline fan speed is used to cross validate with hard-coded baseline fan speed, and a mismatch will waive the threshold passage. For the data section's fan speed signal edge detection, every last 0.1 s signal in a 0.5 s interval after the Header and padding sections is used to perform fan speed classification using signal features (in Table II). Where 0 is represented with base fan speed, and 1 is represented with an elevated fan speed. It is crucial to determine the padding and the following footer section in order to separate the complete data section in the packet. Thus, Data Decoder designs and implements a lookback algorithm. The lookback algorithm starts at the threshold passage timestamp and backtracks the last timestamp that the fan is at base speed. Then, the padding duration is subtracted from that timestamp to obtain the exact timestamp of the last bit. It is also worth mentioning that Data Decoder can read information bit by bit in the SpiralSpy channel using alignment from the header section. Thus, it is possible to perform an early termination, and the already received data is fully decodable. To this end, SpiralSpy communication scheme is complete and successful in transmitting information from air-gapped computing devices to the attacker through the fan-based channel.

VII. SPIRALSPY HIGH-THROUGHPUT OPTIMIZATION

It is worth to mention that current air-gapped computing devices may have multiple fans (e.g., CPU fans, chassis fans, GPU fans, and coolant pump fan) installed to cool different computing components of the system (e.g., CPUs, GPUs, and DRAMs). To further improve the throughput of SpiralSpy, we design an optimization approach that controls multiple fans (i.e., multiple channels) to transmit data simultaneously. Specifically, one central problem is considered before multi-channel information transfer needs to be solved. How to manage all channels to send data? To solve this problem, the fan speed command pipelining algorithm is designed for the new optimization Agent sending process (see Section VII-A, and the Sniper reading process is updated accordingly (see Section VII-B).

A. Agent Sending Process

To utilize multiple fans in a computing device, the header and footer sections of the packet are shared, and the data section is precisely aligned to ensure the integrity of transmission. When multiple fans are present, consecutive commands to control the fan speed can lead to an arbitrary delay in executing each of the succeeding fan's commands. Thus, the multi-channel sending will have a fixed and pre-determined time gap between successive channels without overlapping with existing ones. We design a pipelining algorithm that pads a flexible amount of time between each command to guarantee an even time interval. Without proper padding, Δt_{21} in Fig. 7(a) is significantly longer than other segments, which leads to fan speed overshooting target and never coming back, forfeiting the rest of the data section. Whereas in Fig. 7(b),

each command experiences a flexible period of padding time that guarantees even duration to change fan speed and achieve optimal fan speed sensing. The algorithm used to perform data section modulation is detailed in Alg. 1.

Algorithm 1 SpiralSpy Fan Speed Modulator Algorithm with Multi-Channel Optimization

Input: B : A list of bits
 L_b : Length of list D
 t_{cycle} : Per cycle duration
 t_{acc} : Average acceleration duration
 t_{dec} : Average deceleration duration
 Fan : Fan object to leak data


```

1:  $f_0 = 50$ ;                                ▷ Set baseline fan speed
2: initialize();                              ▷ Gain Fan Control
3: header( $f_0$ );                              ▷ Perform Header Section
4: for  $i_b = 0, 2, \dots, L_b-1$  do
5:    $b = B[i]$ ;                                ▷ Acquire current bit
6:    $s_{exe} = f_0 + b$ ;                          ▷ Compute target fan speed
7:   if  $s_{exe} == Fan.Speed$  then
8:     pause( $t_{cycle}$ );                          ▷ Hold speed for 1 cycle
9:     updateFan( $Fan$ );                          ▷ Update fan object
10:    Continue;                                ▷ Skip following
11:   else if  $b == 1$  then
12:     Accelerate( $t_{acc}$ );                       ▷ Accelerate and Hold
13:     Hold( $t_{cycle} - t_{acc}$ );
14:   else if  $b == 0$  then
15:     Decelerate( $t_{dec}$ );                       ▷ Decelerate and Hold
16:     Hold( $t_{cycle} - t_{dec}$ );
17:   end if
18:   updateFan( $Fan$ );                          ▷ Update fan object
19: end for
20: footer( $f_0$ );                              ▷ Perform Footer Section
21: finalize();                                ▷ Finish the fan control


```

B. Sniper Reading Process

To achieve multiple channel information transfer (e.g., three fans in the target sensing area), we develop a two-layer classification to receive the information.



(a) The initial pipelining without fixed padding.



(b) The new optimized pipelining with fixed padding.

Fig. 7. Agent’s novel padding technique makes fan speed control less volatile, which supports the multi-channel data transfer.

Recognizing the Number of Fans: As the number of fans increases, the quantity of SpiralRF response is also increasing in general. To capture such an increase in SpiralRF response, we utilize spectral and temporal features that effectively describe the signal characteristics. Then, this first-layer model will output the number of fans to the second-layer model.

Classifying Multiple Fans: After getting the number of fans, the second-layer model then uses the feature vector to predict the fan speed for each fan. Due to the different nature of the multi-fan speed control pattern in the data section versus the header/footer section, a special classification scheme is designed. To predict the value for multiple fans in the data section, the number of classes needed to cover all possible combinations in fan speed is humongous (i.e., 76^n for n fans), making it impossible to classify for multiple fans. Thus, a meta-categorization is performed to reduce the number of classes. Given that the fan speeds are mostly changing together, and the maximum difference in fan speed is 1% in the data section, we propose a 2^n class modeling that records all possible combinations directly into the values that the fan speed represents. For instance, a three-fan setup with each fan representing 0 or 1 at any given time, resulting in a total of eight possibilities, thus eight classes in classification. Different from the data section, the header and footer section will require a total of 76 classes in all cases, given multiple fans can reach a threshold speed simultaneously.

Multi-channel Decoding: To append data in the original order for correct decoding, multi-channel decoding implements a metadata mechanism at the beginning of the data section to indicate ordering. For instance, Sniper in a multi-channel (three fans) setup can append three bits in 6 combination orders, and the number of combinations will grow at a factorial rate. Thus, the Agent utilizes first $\log_2(n)$ bits in each channel’s data section to instruct the Sniper the correct ordering, effectively solving the problem while saving space. Upon successful transmission of the data section, file type and extension are obtained using binary signature recognition to further display the file content [41]. To this end, SpiralSpy’s multi-channel optimization can encode, transmit, and decode data through the Spiral-mmWave channel for a stealthy attack.

VIII. SYSTEM IMPLEMENTATION AND EVALUATION SETUP

Fans Preparation: A total of 71 different computing components equipped with cooling fan(s) are used in the following experiments, including 20 GPUs, 20 IoT devices, 20 CPU coolers, and 11 others (e.g., chassis fans). As shown in Fig. 8, a mmWave radar is placed from 25 to 800 cm away from the attacking objects. Representative fans are illustrated, such as NVidia GTX Titan XP GPU cooling fan with 60 mm diameter, CoolerMaster Pro 120 chassis fan with 120 mm diameter, Enermax LIQMAX III 360 CPU AIO fans with 120 mm diameter, and Raspberry Pi case IoT fan with 18 mm diameter [42], [43], [44], [45].

Fan Control Interface: SpiralSpy is capable of stealing data from multiple computing platforms (e.g., PC, IoT/Edge devices, cloud server) deployed in a variety of application scenarios. In our implementation, we use various computing component drivers to accurately control the fans’ speed without

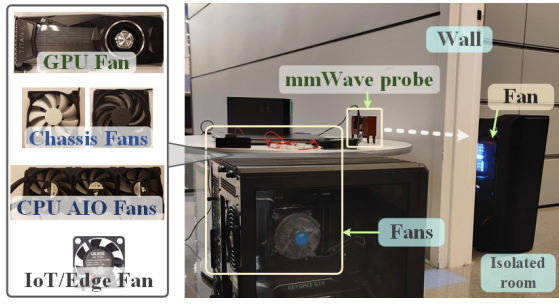


Fig. 8. The setup for the SpiralSpy evaluation mainly consists of the mmWave probe, cooling fans, and the isolated room.

triggering system or user alerts (e.g., "nvidia-smi" command for NVidia's GPU fans, Windows Management Instrumentation for CPU fans, PWM pin control for IoT fans, and IPMItool for cloud server fans) [46], [47], [48], [49], [31], [50].

mmWave Probe: The mmWave-Spiral sensing device implementation is a 24 GHz FMCW radar that consists of a pair of 4×4 antenna array for each of TX and RX terminals as well as 450 MHz bandwidth [51], which offers an antenna directionality of 19.8 dBi. The weight of the device is 45.5g, and the total manufacturing cost is less than \$100.

Data Collection: For performance analysis, SpiralRF response samples are collected from various fans and multiple sensing parameters. The default sensing distance is 1 meter. Each fan is programmed to stay at a preset fan speed (e.g., 50% PWM duty cycle/2400 rpm) with continuous mmWave samples under correct speed labeling. For instance, a fan with a variable speed range from 30% duty cycle to 100% duty cycle will result in 71 different speeds, and each speed will be recorded for at least 1 second. This process then iterates for 100 cycles, resulting in 7,100 traces in each trail. In total, 106,500 traces are recorded for training, and 35,500 traces are recorded for testing.

Evaluation Metrics: In order to examine the performance of SpiralSpy, the confusion matrix [52] is adopted for better visualization. The attack result (i.e., predicted fan speed) versus ground truth fan speed in the percentage of maximum fan speed will be directly compared. In addition, metrics such as attack accuracy, precision, recall, and correlation coefficient [53] are used for a more comprehensive analysis.

IX. PERFORMANCE EVALUATION

In this section, we analyze SpiralSpy's performance across multiple fans under different user scenarios. When experimented against GPU, CPU, and IoT fans, SpiralSpy shows strong performance as following.

A. Performance on GPUs

General-purpose GPUs have become a more and more significant module in modern computing devices due to their parallel computing power accelerated 3D graphics processing as well as machine learning. GPUs targeting different customer markets may be designed different (e.g., from one turbofan to three multi-blade fans) cooling solution very severely. Thus, we first examine SpiralSpy's attack performance across 20 different GPUs equipped with cooling fan(s). Each of the GPUs is plugged into the motherboard's top PCI Express slot to maximize the utilization of 16 lanes bandwidth, which prevents

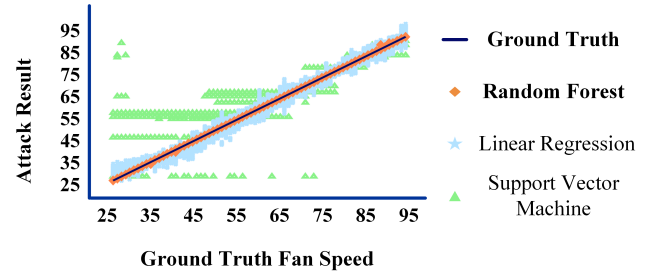


Fig. 9. SpiralSpy attack performance with three different prediction mechanisms on GPUs. The random forest model shows the superior performance.

GPU processing performance in theory. In all scenarios, the GPU fans are set to face downward (i.e., pushing air from the bottom of the chassis to GPU heat sinks) and run at different fan speeds at 1% speed increase interval. The attack results show exceptional performance and reach a 99.99% average attack accuracy where the leaked bits are exactly the same as the ground truth. To better illustrate it, attack result fan speed versus ground truth fan speed is plotted in Fig. 9. The high accuracy is also coupled with a correlation coefficient of 1, which means the prediction is exceptionally close to the ground truth, and the errors in the sample are insignificant. The precision and recall are both 99.89%, implying the features in the spectral and temporal domain accurately describe the SpiralRF response indeed instead of the background noises. Moreover, compared to the accuracy of 7% and 17% obtained by the support vector machine and linear regression, respectively, SpiralSpy's random forest model demonstrated superiority in attack success rate.

B. Performance on IoT Devices

More and more IoT/edge devices have powerful computing resources and are deployed as computing nodes. Typically, IoT devices are designed with wireless modules and built-in antennas and thus enclosed in a plastic casing with an opening to allow air cooling. To comprehensively evaluate the performance of these fans, this experiment tests with 20 different IoT devices equipped with cooling fans. To avoid air flow scarcity at low speed on small fans, the fan speed attack is set from 50% to 100%, resulting in 51 total classification categories. As shown in Fig. 10, SpiralSpy is able to reach the average accuracy at 92.79% with both precision and recall being 92.75%. The random forest model displayed superior performance compared to the linear and support vector models. Since IoT fans are smaller in size (<30 mm diameter), which leads to a smaller radar cross-section, the attack performance is slightly lower than the GPU attack results. In general, results demonstrate SpiralSpy's capacity to attack IoT/edge devices even with small form-factor fans.

C. Performance on CPU Coolings

The structure of the fan on CPU cooling is different from the previous two types. Given the properties of a working CPU, low fan speed and high airflow fans are designed with unique elevated blade shape [54], such fan with a maximum fan speed of 1200 RPM which means only 12 RPM difference per 1% duty cycle change. This challenges SpiralSpy to attack these fans with significant fine-grained granularity. This experiment tests with 20 different off-the-shelf cooling solutions for a

variety of CPUs. As shown in Fig. 11, despite the granularity challenge, *SpiralSpy* reached 99.9% average attack accuracy in 71 class fan speed classification (30% to 100% duty cycle) and a small portion of error near 30% duty cycle where the fan itself was not stable enough to maintain speed. These results suggest that *SpiralSpy* model is accurate at attacking multiple designs of CPU fans.

D. Performance on Multi-Channel Attacks

To assess multi-channel *SpiralSpy* attack performance, we adopt a confusion matrix to show attack result versus ground-truth value in a multi-channel data leakage (three fans simultaneously). This experiment tests with five different three-fans sets. Overall, as shown in Fig. 12, the average attack result has a very close relationship with the ground truth and yields a correlation coefficient of 0.9985, accompanied by accuracy, precision, and recall all being 99.95%. Besides, *SpiralSpy* can achieve up to 6 bps throughput rate due to 2 bps per fan simultaneously. It is also worth mentioning that the attack result shows a small fraction of outlining errors, which is due to unexpected gains in the amplitude of *SpiralRF* modulated frequencies that lead to far-off classification. Although such outlining factors are challenging to model against, it proves that the non-linearity in the *SpiralRF* response was effectively captured by our model and works for the multi-channel attack.

X. ROBUSTNESS AND PRACTICALITY EVALUATION

We evaluate the robustness and practicality of *SpiralSpy* under various real-world circumstances. For each test, we employ 25 different computing devices with fan(s) (unless specifically mentioned) and follow the same implementation mentioned in Section VIII.




Fig. 10. *SpiralSpy* attack performance on IoT devices.




Fig. 11. *SpiralSpy* attack performance on CPUs. Random forest model is again superior.




Fig. 12. *SpiralSpy* multi-channel attack performance with a three-fans set. Each fan’s attack result is represented by either 0 or 1 in the value label.

A. Impact of Attack Orientation

To better understand the importance of tangential velocity sensing introduced in Sec. III-B, we evaluate the velocity type. In each test, the probe orientation changes from -90° to 90° , and the average attack accuracy are shown in Fig. 13. The attack performance of *SpiralSpy* peaks at 0° sensing angle (i.e., probe directly facing the fan in upright position), which is based on **the tangential velocity** of the fan. Besides, the system maintains relatively high performance as the mmWave probe moves $\pm 30^\circ$ to the side. However, once the probe moves beyond $\pm 60^\circ$ from the center (i.e., sensing on the radial velocity of the fan), the accuracy drops under 80%. As a result from *SpiralSpy*’s robustness design in reading fan speed, it excels at attacking with different angles with strong performance, which also proves that sensing tangential velocity is superior to the radial velocity in this covert channel attack.




Fig. 13. Performance impact of different velocity types.




Fig. 14. Attack performance under the influence of the sensing distance.

B. Impact of Attack Distances

Attack distance is a crucial factor in covert channel data leakage that determines possible attack scenarios to the adversary. We, therefore, examine the impact of the attack distance from 0.25 m to 8 m, considering both mmWave reachable distance and attenuation of *SpiralRF* response over the long range. As shown in Fig. 14, *SpiralSpy* maintains the attack accuracy to above 99% even at the distance of eight meters, enabling a feasible covert attack without proximity requirement. Considering the model was trained on 1 m, the robustness of *SpiralSpy* against distance scale is excellent. Compared to existing covert channels’ maximum distance of 1.6 m as shown in Table III, *SpiralSpy*’s high performance at eight meters expands new possibilities.

C. Impact of Computing Workloads

Computing workloads on devices may undermine the fans’ stability while spinning at a high rate and eventually result in fluctuations of fan speeds. Therefore, it is crucial to examine *SpiralSpy*’s performance under various fan workloads. Specifically, in our experiment, the workloads of CPU and

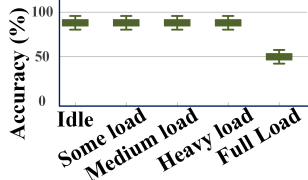


Fig. 15. Attack performance with different computing workloads.

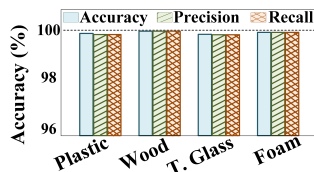


Fig. 16. Evaluation to determine the influence of isolation material.

GPU are manipulated by running different tasks reaching 20%, 40%, 60%, 80%, 100% usage time. As shown in Fig. 15, the attack accuracy of *SpiralSpy* is consistently above 99% up to the extreme load usage where accelerations are utterly slow. Overall, *SpiralSpy* is robust and capable of running at various computer workloads with a strong performance.

D. Impact of Through-Wall Isolations

One of the key features of *SpiralSpy* is being a through-wall covert channel attack. In real-world conditions, different materials of occlusion will be present. Thus, it is crucial to examine *SpiralSpy*'s capability to attack fan with different occlusion scenarios. To simulate various application scenarios, popular materials such as plastic, wood, tempered glass (i.e., T. glass), and foam are set between the computing device and the Sniper. As shown in Fig. 16, the performance of *SpiralSpy* reaches above 99% with all four occlusion scenarios, proving *SpiralSpy*'s robustness against fan isolation.

E. Impact of Ambient Dynamics

Another essential consideration in attack is ambient interference. We select four common influential factors on our attack system in our daily life, including magnetism, vibration, and human interference, and air humidity. (1) high humidity scenario where humidifier directly pumps moisture in front of the fan; (2) a high magnetic field scenario where a one cm diameter grade Y30 magnet [55] is placed in front of the fan, one cm away from the rotating center; (3) a high vibration scenario where a smartphone is placed on the chassis, introducing constant vibration; (4) and a human interference scenario where a subject is walking in the radar line of sight but not blocking the fan. As shown in Fig. 17, *SpiralSpy* maintains high attack accuracy despite substantial interference such as moisturized air being pulled into the fan. Given *SpiralSpy*'s system design filtering out nonrelated information in RF signals, *SpiralSpy* is robust against environmental interference with attack accuracy above 98% in all four scenarios.

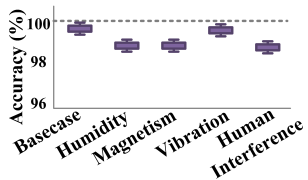


Fig. 17. Performance impact of different environmental factors.

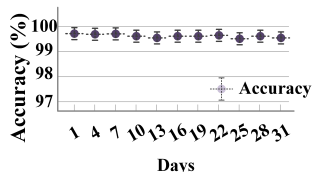


Fig. 18. Attach performance against one-month time change.

F. Sustained Attack Analysis

The durability of a covert channel attack is crucial in information theft. To prove the long permanence of *SpiralSpy*, a month-long study is performed to understand the performance stability over time. At the interval of 3 days, we collect a new data sample to test the pre-trained model. As illustrated in Fig. 18, the sensing accuracy in this durability test never fell below 99%, indicating *SpiralSpy*'s robustness within an extended period.

	Text File (Password)	Audio File (“Okay”)	Image File (Iris Code)
Ground Truth	A1b2!@/		
Single-C Result	A1b2!@/		
Multi-C Result	A1b2!@/		

Fig. 19. Single-channel and multi-channel attacks on three different types of private data files (i.e., text, audio and image).

G. Multi-channel Attack Analysis

To examine *SpiralSpy*'s capability in stealthily leaking data in multi-channel scenarios, we design experiments to transmit different types of privacy-related data (e.g., the password for financial accounts [56], iris code image that unlocks smartphone/door-lock [57], and confirmation voice audio that banks often use for user authentication in phone calls [58]). In each of the three trails, the attack based on a single fan with Corsair ML140 (Single-C) and multi-channel fans (three fans) from Enermax CPU AIO (Multi-C) are conducted, which test with text password (i.e., “A1b2!@/” in plain text .TXT file - 42 bits), voice confirmation (i.e., recorded audio data “Okay” in .MP3 file - 5400 bits), and partial iris code image (i.e., 60 x 60 black and white image in binary logical representation - 3600 bits). For the single-channel attack, the transmission duration of the three trails lasts 1, 46, and 31 minutes with an average accuracy of 99.9%, 99.4%, 99.1%, respectively. The transmission duration for three attacks verifies 6 bps rate deducted from 2 bps per fan. For the multi-channel attack, the transmission duration is 1, 16, and 11 minutes with an average accuracy of 99.6%, 99.1%, 99.4%, respectively. Fig. 19 shows the three attack trails, including their ground truth and attack results. Clearly, as we can see, the reconstructed text, audio, and image are almost identical to the ground truth, showcasing the severe security implication of *SpiralSpy* in a multi-channel attack.

XI. REAL-WORLD ATTACK

Experimental Setup: Due to the accessibility and portability of the setup, the attack system can access the target fan used in various public spaces. Thus, to examine *SpiralSpy*'s attack performance in the real world, we conduct the real-world fan-based covert channel attack in two scenarios. In the first scenario, the air-gap prevents the attacker from being in close proximity with the computing device. The mmWave probing

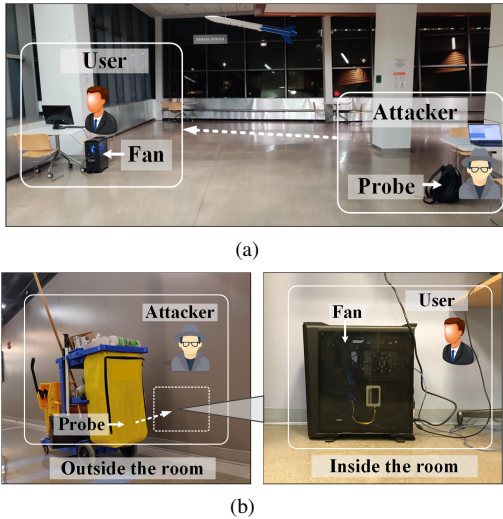


Fig. 20. Real-world attacks are conducted in two everyday scenarios: (a) Scenario 1: The attacker infers the credential information from a remote location in a hall; (b) Scenario 2: The attacker implements a through-wall inference against an isolated room. The attacker can pretend as a staff, and the probe can be hidden in a normal handbag or cleaning cart, arousing no suspicion to the victim and nearby people.

device is hidden in a carrier container that does not raise suspicion (e.g., backpack, laptop sleeve bag). The attack distance is set to be 8m. Next, in the through-wall scenario, a thick wall (i.e., 20 cm) is employed to test SpiralSpy’s through-wall attack capabilities. Furthermore, the air-gap forces the attacker to disguise/fit into an object that would naturally appear to be stationary on the other side of the wall (e.g., cleaning cart in office buildings). The attack distance is 70cm. As shown in Fig. 20, the mmWave probing devices are placed air-gapped from the victim computer. It is worth mentioning that each scenario will use three types of data covering three common file media (text, audio, and image), and each file type will be attacked using both the single-channel and multi-channel methods.

Evaluation Results: As shown in Fig. 21, the attack results for three file types in two scenarios are illustrated. The attacks for meeting detail text (240 bits), token detail text (272 bits), fingerprint image (3600 bits), and confirmation audio file (4200 bits) took 3, 3, 31, and 36 minutes in a single-channel (Single-C) attack, and 2, 2, 11, and 13 minutes in multi-channel (Multi-C), respectively. Given the fixed time interval in protocol packet design, no bits are ever lost or misplaced in the attack results. The performance of SpiralSpy achieves an average of 99.4 % across text file experiments. Some bits are misinterpreted by the Sniper resulting. However, these typo-like words in text files produce little impact on information theft. For instance, the letter “g” in word “building” represents 103 decimal or 1100111 binary in 8 bit ASCII encoding, as the fourth bit 0 is mistaken with 1, the decimal value is changed to 111, resulting in the letter “o” and the word “buildino” in Multi-C S2. Besides, for the audio file and the fingerprint image, the average accuracy is 99.4% and 99.1%, respectively. Some binary errors can be discovered in results, yet these look like slight “noise” and do not affect the main characteristics of audio and fingerprint. In the audio file, a binary error does minor damage to overall sound since audio encoding such as MP3 segments audio into different parts [59]. For instance, a part of the data stream “...1001100...” is incorrectly recognized

as “...1001100...”. Also, in the fingerprinting image, a binary error flips a black/white representation of a pixel, changing the little ridge’s pattern. For example, one part, “...00010100...” is mistakenly intercepted as “...00011100...”. Moreover, generally in the attack results, the audio sounds nearly identical to the ground truth, and the fingerprint is indistinguishable from the ground truth image by humans and machines (e.g., the speaker identification [60], and fingerprint minutiae matching [61]). Thus, these results suggest that this attack system provides reliable performance in actual practice.

	Text File (Meeting Detail)	Text File (Token Detail)	Audio File (“I Consent”)	Image File (Fingerprint)
Ground Truth	Meet on 6 PM at Building X for weapons	1JMK91gy6MUBuySoxoArB6MtyeNhhSa7dr		
Single-C Result (S1)	Meet on 6 PM at Building X for weapons	1JMK91gy6MUBuySoxoArB6MtyeNhhSa7dr		
Multi-C Result (S1)	Meet on 6 PM at Building X for weapons	1JMK91gy6MUBuySoxoArB6MtyeNhhSa6dr		
Single-C Result (S2)	Meet on 6 PM at Building X for weapons	1JMK91gy6MUBuySoxoArB6MtyeNhhSa7dr		
Multi-C Result (S2)	Meet on 6 PM at Buildino X for weapons	1JMK91gy6MUBuySoxoArB6MtyeNhhSa7dr		

Fig. 21. Real-world attack results in two different scenarios against the ground truth. (S) is short for Scenario.

XII. COUNTERMEASURE ANALYSIS

In this section, we will discuss two sets of practical countermeasures against SpiralSpy in the following section. The first countermeasure set (i-iv) is the physical-based protection, altering the working environment, user behavior, or hardware to mitigate the security risk. The second set (v-vi) is the virtual-based countermeasure, a purely software-based solution with no hardware or working procedure requirement.

(i) Large Isolation Zone: One of the most intuitive defenses is to create a substantially large isolation zone (e.g., thousands of square feet), which is considered effective against most of the air-gapped attacks, including SpiralSpy. However, it is not practical (e.g., due to cost and usability) in real-world scenarios.

(ii) Shielding the Device: Electromagnetic Field (EMF) shielding (e.g., aluminum foil) can be applied to the computing devices with fans to block the mmWave sensing signal, as shown in Fig. 22(a). The shielding on the device can eliminate mmWave from going in or out the device to prevent SpiralSpy covert channel attack. However, there are few practical concerns when implementing this method on computing devices. First, using EMF shielding material to cover the device case/fan set will inevitably cause thermal regulation issues due to lack of heat dissipation, making the device difficult to operate or inoperable. Besides, deploying the shielding material increases the extra cost. In addition, another option is to employ a professional electromagnetic shielded room for isolation. However, this solution will substantially increase the cost and is difficult to scale.

(iii) **RF Interference:** Another possible way to eliminate this covert channel is to block the mmWave receiving channel with a radio frequency interfering device (e.g., mmWave jammer), as shown in Fig. 22(b). By doing so, mmWave receiving terminals in the surrounding will be flooded with full amplitude waves in the spectrum and fail to perform fan speed recognition. However, it is challenging to implement without knowing the *SpiralSpy* carrier mmWave frequency in actual practice. Besides, this jammer device can harm the health of the staff nearby when working for a long time.

(iv) **Honeypot Device/Fan:** It is also possible to prevent this covert channel by deploying a similar extra computing device/fan nearby to spoof and deflect the mmWave sniper. However, *SpiralSpy* can resist this countermeasure with a specially designated header and footer section of the packet. Besides, benefited from the EWT signal separation, *SpiralSpy* can eradicate the noise from the extra fan at the initial step in the reading scheme. Moreover, it is arduous to figure out the *SpiralSpy* information transmission protocol coincidentally in actual practice. Also, it will increase the protection cost.

(v) **Blocking the Fan Control Interface:** Since the covert channel is based on the fan status, the system can completely lock the fan control interface to disable *SpiralSpy*'s ability to transmit data. However, this solution leads to the sizeable irritating noise or supernumerary power consumption from the fan, which can heavily affect the ordinary usage of the authenticated users when they have no fan speed control access. Moreover, many professional or commercial software utilize fine-grained fan control to keep the high performance [29], [30], [31]. If the device system blocks the fan control interface, it can also impair these software performances and user working efficiency.

(vi) **Anti-Agent Scan:** In highly secure computing scenarios such as air-gapped computing, it is natural to install some form of anti-virus software to routinely check malicious activities (e.g., abnormal network traffic). Given *SpiralSpy*'s covert channel does not utilize any network module (e.g., Ethernet, Wi-Fi, Bluetooth), it is hard for existing anti-virus software to recognize *SpiralSpy*'s program as malicious through network traffic monitoring. Furthermore, *SpiralSpy*'s software program utilizes driver-level command provided by hardware manufacturers that manipulates fan speed as if the user is acting on the switches. Such design tricks the user by acting as OS thermal regulatory operation and tricks anti-virus by pretending as user actions. During the evaluation, top popular anti-virus protections (all latest versions), Windows 10 Defender, HuoRong Security, and McAfee Security Scan, do not report any security issues and warnings about this *SpiralSpy*'s Agent. Additionally, it is possible to design specific anti-virus protections for fan modulation. However, this solution is still considered inadequate since it is hard to differentiate the fan modulation resulting from the Agent or other professional or commercial software.

XIII. DISCUSSION

SpiralSpy Stealthiness Analysis: *SpiralSpy* is designed to maintain stealthiness while transmitting data through fan speed control. While transmitting the header and footer section, the noise increase measured from a turbo GPU fan from idle to




Fig. 22. Examples of countermeasures for *SpiralSpy*: (a) shielding the device case with aluminum foil; (b) the covert channel interference with an RF jamming device.

75% fan speed is +2 to +4 dB at 10 centimeter distance, where the ambient noise is at 34 dB. Compared to the background noise, it is hardly distinguishable by the human ear at a 100 cm distance. For the data section, because *SpiralSpy*'s acceleration and deceleration is kept at minimal (i.e., $\pm 1\%$ of maximum fan speed), the fan speed change is unnoticeable acoustically. Since *SpiralSpy* does not rely on traditional networking infrastructure to transmit data. Moreover, *SpiralSpy* causes little to no variation in power consumption of the fan (less than 0.01%), making it stealthy toward power monitoring applications as well.

Fan Speed Adjustment: In order to further increase the data transmission rate, it is possible to increase acceleration and deceleration curves through more fine-tuned control (i.e., shorten the time to change speed). For example, it is an option to improve the acceleration speed by modifying the current hardware and applying a higher voltage on the fan. However, such a solution is not suitable for this attack and inevitably makes notable noise.

Strict Scrutiny Against the Agent Implantation: In order to counter *SpiralSpy*'s Agent implantation through social engineering or supply chain attack, security protocols for access control can be more strictly implemented and enforced. Sadly, catching a security breach through social engineering is still a complex and classic security problem involving devices to humans, which remains a cat and mouse race that is exceptionally challenging to patch thoroughly.

Attack Distance: The attack distance can be potentially prolonged to more than 10 m with a set of professional instruments in mmWave band toward air-gapped computing devices [63]. However, this solution dramatically increases the attack cost (5X higher cost).

Agent Optimization: Representative information encoding methods, such as SHA-256 [64], can be utilized for the Agent sending process. However, such optimization requires more computation resources, easier to trigger alerts in usage monitoring. Moreover, the information encoding method selection is out of the scope of this attack work. Besides, it is an option to divide the packet data into multiple smaller packets and send them sequentially to increase the attack robustness. However, it inevitably costs more time in information transmission, sacrificing the attack speed and stealthiness.

Future Applications: This covert channel opens up possibilities for various application scenarios that involve both RF sensing and rotating objects. Considering the ubiquity of objects that carries rotational motion in the real world (e.g., vehicle tires, motors, drone propellers), this work has

TABLE III. A COMPARISON OF COVERT CHANNELS ON AIR-GAPPED COMPUTING DEVICE

System	Channel Modalities	Distance	Stealthiness	Bit rate	Through-Wall	Multi-Channel
BitWhisper [10]	Thermal	0.4 m	No (CPU Full Load)	$\ll 1$ bps	No	No
GSMem [20]	EM Leakage	1 m	No (Memory Full Load)	2 bps	No	No
Odini [11]	Magnetic	1.5 m	No (CPU Full Load)	1 bps	Yes	Yes
AiR-ViBeR [62]	Vibration	1.6 m	No (Loud Noise from max fan speed)	0.5 bps	No	No
Fansmitter [19]	Acoustic	1 m	No (Loud Noise from max fan speed)	0.25 bps	No	No
SpiralSpy	Tangential Velocity	8 m	Yes	2-6 bps	Yes	Yes

the potential to turn objects carrying fine and rapid rotational motion control into a SpiralSpy channel.

XIV. RELATED WORK

Air-Gapped Computing Device Covert Channels: Several attempts have been made in the literature to explore the covert channels on air-gapped computing devices. Such attacks usually leverage nontraditional communication media to enable data exfiltration or collude, such as acoustic [65], [19], [66], optic [67], [68], RF [20], [69], [70], thermal [10], [71], and magnetic [11] artifacts. However, covert channels based on these modalities either demand proximity or require network access, leaving mitigation against them easy to find. Meanwhile, their communication capability is very limited that impractically poses threats in real practice. Using a low-cost mmWave radar, SpiralSpy is the first work that utilizes the SpiralRF response of rotational fans to realize a covert channel, while having minimal thermal, acoustic, and computation usage impact on normal computation operation. As shown in Table III, SpiralSpy demonstrates substantial superiority in long attack distance, fast data transmission, through-wall, and multi-channel as a stealthy and practical covert channel.

mmWave Sensing: Originally invented as a means for wireless communication (e.g., 5G networking) [72], [73], recent years are witnessing a growing interest in using mmWave radios for mobile sensing. Owing to its fine-grained ranging ability, mmWave signals are found effective for location-based services such as navigation, mapping, and object reconstruction/imaging [74], [75], [76], [77], [78]. In the regime of the mmWave radar, researchers utilize the unique relationship between the signal reflection characteristics and objects in the field of view for identifying material types and object status (e.g., pattern, temperature, metallicity and wear etc.) [79], [80], [81], [82], [83], [84]. A particular property of mmWave signals related to our work lies in their micro-Doppler effect (c.f. Section. III) or the power distribution, by which one can estimate an object’s micro-motion for a range of applications, such as vital sign monitoring [85], vibration detection [86], [87], [88], gesture recognition [89], and hand gesture tracking [90]. However, in contrast to the prior arts that are based on the radial velocity of a moving object, SpiralSpy is the first to utilize *the tangential velocity* of the fan for the covert channel attack and thereby provides more sensing granularity than the classic micro-Doppler effect in this work.

XV. CONCLUSION

This paper identified and designed a stealthy and practical fan-based covert channel attack on air-gapped computing devices using mmWave sensing. SpiralSpy exploited the SpiralRF response based on the tangential velocity of the fan. We designed the customized information encoder and fan speed modulator on the victim device to send the credentials. Then, we developed an end-to-end SpiralSpy sniper to extract the fan status and decode the credential information after receiving the SpiralRF response from the fans. Finally, we optimized the multi-channel attack with multiply fans to magnify the information transmission (one-way communication) throughput. SpiralSpy achieved information theft with a rate of 6 bits per second in 8 meters without computer hardware and driver modification. It recommends sensitive and protected computing devices to be careful about this new covert channel attack and provides a new view further to explore the security of the air-gapped computing devices.

ACKNOWLEDGMENTS

We thank our shepherd Dr. Ivan De Oliveira Nunes and all anonymous reviewers for their insightful comments on this work. This work was in part supported by the National Science Foundation under grant No. CNS-1718375, ECCS-2028872 and CNS-2050910.

REFERENCES

- [1] “Billions of password leaked,” Jan 2021. [Online]. Available: <https://www.wired.com/story/collection-leak-username-passwords-billions/>
- [2] B. W. Lampson, “A note on the confinement problem,” *Communications of the ACM*, vol. 16, no. 10, pp. 613–615, 1973.
- [3] S. J. Murdoch and S. Lewis, “Embedding covert channels into tcp/ip,” in *International Workshop on Information Hiding*. Springer, 2005, pp. 247–261.
- [4] J. Xing, A. Morrison, and A. Chen, “Netwarden: Mitigating network covert channels without performance loss,” in *11th {USENIX} Workshop on Hot Topics in Cloud Computing (HotCloud 19)*, 2019.
- [5] Y. Yarom and K. Falkner, “Flush+ reload: a high resolution, low noise, L3 cache side-channel attack,” in *23rd {USENIX} Security Symposium ({USENIX} Security 14)*, 2014, pp. 719–732.
- [6] W.-M. Hu, “Lattice scheduling and covert channels,” in *Proceedings 1992 IEEE Computer Society Symposium on Research in Security and Privacy*. IEEE Computer Society, 1992, pp. 52–52.
- [7] M. Guri, M. Monitz, and Y. Elovici, “Usbee: Air-gap covert-channel via electromagnetic emission from usb,” in *2016 14th Annual Conference on Privacy, Security and Trust (PST)*. IEEE, 2016, pp. 264–268.

- [8] Z. Wang and R. B. Lee, "Covert and side channels due to processor architecture," in 2006 22nd Annual Computer Security Applications Conference (ACSAC'06). IEEE, 2006, pp. 473–482.
- [9] "To air-gap or not air-gap industrial control networks," Jan 2021. [Online]. Available: <https://www.tripwire.com/state-of-security/ics-security/air-gap-industrial-control-networks/>
- [10] M. Guri, M. Monitz, Y. Mirski, and Y. Elovici, "Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations," in 2015 IEEE 28th Computer Security Foundations Symposium. IEEE, 2015, pp. 276–289.
- [11] M. Guri, B. Zadov, and Y. Elovici, "Odini: Escaping sensitive data from faraday-caged, air-gapped computers via magnetic fields," IEEE Transactions on Information Forensics and Security, vol. 15, pp. 1190–1203, 2019.
- [12] 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 Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services, 2020, pp. 148–160.
- [13] A. Wang, Z. Li, C. Peng, G. Shen, G. Fang, and B. Zeng, "Inframe++ achieve simultaneous screen-human viewing and hidden screen-camera communication," in Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services, 2015, pp. 181–195.
- [14] P. K. Kerr, J. Rollins, and C. A. Theohary, The stuxnet computer worm: Harbinger of an emerging warfare capability. Congressional Research Service Washington, DC, 2010.
- [15] "Military computer attack confirmed," Jan 2021. [Online]. Available: <https://www.nytimes.com/2010/08/26/technology/26cyber.html>
- [16] "Department of homeland security, ics-cert monitor, malware infections in the control environment," Jan 2021. [Online]. Available: <https://archive.org/details/5981914-National-Security-Archive-Department-of-Homeland>
- [17] "Social engineering, the usb way," Jan 2021. [Online]. Available: <https://www.darkreading.com/attacks-breaches/social-engineering-the-usb-way/d/d-id/1128081?>
- [18] "Supply chain attacks are congruent with insider threat," Jan 2021. [Online]. Available: <https://dzone.com/articles/cyber-supply-chain-risk>
- [19] M. Guri, Y. Solewicz, A. Daidakulov, and Y. Elovici, "Fansmitter: Acoustic data exfiltration from (speakerless) air-gapped computers," arXiv preprint arXiv:1606.05915, 2016.
- [20] 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), 2015, pp. 849–864.
- [21] W.-C. Chen and Y.-Y. Tzou, "Current-mode sensorless control of single-phase brushless dc fan motors," in 2011 IEEE Ninth International Conference on Power Electronics and Drive Systems. IEEE, 2011, pp. 659–663.
- [22] A. Arredondo, P. Roy, and E. Wofford, "Implementing pwm fan speed control within a computer chassis power supply," in Twentieth Annual IEEE Applied Power Electronics Conference and Exposition, 2005. APEC 2005., vol. 1. IEEE, 2005, pp. 148–151.
- [23] "Cooling computers," Feb 2021. [Online]. Available: <https://www.intel.com/content/www/us/en/gaming/resources/cpu-cooler-liquid-cooling-vs-air-cooling.html>
- [24] R. Lanzafame and M. Messina, "Fluid dynamics wind turbine design: Critical analysis, optimization and application of bem theory," Renewable energy, vol. 32, no. 14, pp. 2291–2305, 2007.
- [25] C. Wu, F. Zhang, Y. Fan, and K. R. Liu, "Rf-based inertial measurement," in Proceedings of the ACM Special Interest Group on Data Communication, 2019, pp. 117–129.
- [26] F. Xu, T. Hong, J. Zhao, and T. Yang, "Detection and identification technology of rotor unmanned aerial vehicles in 5g scene," International Journal of Distributed Sensor Networks, vol. 15, no. 6, p. 1550147719853990, 2019.
- [27] V. C. Chen, F. Li, S.-S. Ho, and H. Wechsler, "Micro-doppler effect in radar: phenomenon, model, and simulation study," IEEE Transactions on Aerospace and electronic systems, vol. 42, no. 1, pp. 2–21, 2006.
- [28] Z. Peng, L. Ran, and C. Li, "A 24-ghz low-cost continuous beam steering phased array for indoor smart radar," in 2015 IEEE 58th International Midwest Symposium on Circuits and Systems (MWSCAS). IEEE, 2015, pp. 1–4.
- [29] "Msi afterburner," Feb 2021. [Online]. Available: <https://www.msi.com/Landing/afterburner/graphics-cards>
- [30] "Asus ai cooling," March 2021. [Online]. Available: <https://www.asus.com/microsite/motherboard/Intelligent-motherboard/AI-Cooling.html/>
- [31] "Nvidia fan control," Dec 2020. [Online]. Available: <https://developer.download.nvidia.com/compute/DCGM/docs/nvidia-smi-367.38.pdf/>
- [32] "Corsair hydro series h100i v2 extreme performance water, liquid cpu cooler," Jan 2021. [Online]. Available: <https://dzone.com/articles/cyber-supply-chain-risk>
- [33] R. De Buda, "Coherent demodulation of frequency-shift keying with low deviation ratio," IEEE transactions on communications, vol. 20, no. 3, pp. 429–435, 1972.
- [34] F. Xiong, "M-ary amplitude shift keying ofdm system," IEEE Transactions on Communications, vol. 51, no. 10, pp. 1638–1642, 2003.
- [35] R. v. Nee and R. Prasad, OFDM for wireless multimedia communications. Artech House, Inc., 2000.
- [36] F. Zheng, G. Zhang, and Z. Song, "Comparison of different implementations of mfcc," Journal of Computer science and Technology, vol. 16, no. 6, pp. 582–589, 2001.
- [37] J. Gilles, "Empirical wavelet transform," IEEE transactions on signal processing, vol. 61, no. 16, pp. 3999–4010, 2013.
- [38] S. Weisberg, Applied linear regression. John Wiley & Sons, 2005, vol. 528.
- [39] J. A. Suykens and J. Vandewalle, "Least squares support vector machine classifiers," Neural processing letters, vol. 9, no. 3, pp. 293–300, 1999.
- [40] V. Svetnik, A. Liaw, C. Tong, J. C. Culberson, R. P. Sheridan, and B. P. Feuston, "Random forest: a classification and regression tool for compound classification and qsar modeling," Journal of chemical information and computer sciences, vol. 43, no. 6, pp. 1947–1958, 2003.
- [41] "Trid - file identifier," Feb 2021. [Online]. Available: <https://mark0.net/soft-trid-e.html>
- [42] "Nvidia titan xp," Jan 2021. [Online]. Available: <https://www.nvidia.com/en-us/titan/titan-xp/>
- [43] "Raspberry pi 4 case fan," Jan 2021. [Online]. Available: <https://www.raspberrypi.org/products/raspberry-pi-4-case-fan/>
- [44] "Cooler master pro 120 af," Jan 2021. [Online]. Available: <https://www.coolermaster.com/catalog/coolers/case-fan/masterfan-pro-120-af/>
- [45] "Enermax iii 360," Jan 2021. [Online]. Available: <https://www.enermax.com/en/products/liqmax-iii-360-argb>
- [46] "Hp server fan controller," Dec 2020. [Online]. Available: <https://github.com/damianmoore/hp-server-fan-controller>
- [47] "ipmitool," Dec 2020. [Online]. Available: <https://github.com/ipmitool/ipmitool>
- [48] "Smc mac osx fan control," Dec 2020. [Online]. Available: <https://github.com/hholtmann/smcFanControl/>
- [49] "Windows speedfan fan control," Dec 2020. [Online]. Available: <http://www.almico.com/speedfan.php/>
- [50] "Raspberry pi pwm," Jan 2021. [Online]. Available: <https://www.mbtechworks.com/projects/raspberry-pi-pwm.html/>
- [51] Z. Peng, J.-M. Muñoz-Ferreras, R. Gómez-García, L. Ran, and C. Li, "24-ghz biomedical radar on flexible substrate for isar imaging," in 2016 IEEE MTT-S International Wireless Symposium (IWS). IEEE, 2016, pp. 1–4.
- [52] J. T. Townsend, "Theoretical analysis of an alphabetic confusion matrix," Perception & Psychophysics, vol. 9, no. 1, pp. 40–50, 1971.
- [53] R. A. Fisher, "Statistical methods for research workers," in Breakthroughs in statistics. Springer, 1992, pp. 66–70.
- [54] "Noctua fan," Jan 2021. [Online]. Available: <https://noctua.at/en/nfs12b-redux-1200-pwm/>
- [55] "Magnet grades," Feb 2021. [Online]. Available: <https://www.first4magnets.com/tech-centre-i61/information-and-articles-i70/ferrite-magnet-information-i83/grades-of-ferrite-magnets-i106>

- [56] H. Gilbert and H. Handschuh, "Security analysis of sha-256 and sisters," in *International workshop on selected areas in cryptography*. Springer, 2003, pp. 175–193.
- [57] J. Daugman, "Information theory and the iricode," *IEEE transactions on information forensics and security*, vol. 11, no. 2, pp. 400–409, 2015.
- [58] P. R. Kennedy, T. G. Hall, and W. C. Yip, "Radio telecommunication device and method of authenticating a user with a voice authentication token," Jul. 4 2000, uS Patent 6,084,967.
- [59] J. Sterne, *MP3: The meaning of a format*. Duke University Press, 2012.
- [60] "speech signal based speaker identification," Accessed: 2021-5-1. [Online]. Available: https://www.mathworks.com/matlabcentral/fileexchange/68230-speech-signal-based-speaker-identification?s_tid=srchtitle
- [61] "Fingerprint minutiae extraction," Accessed: 2021-5-1. [Online]. Available: https://www.mathworks.com/matlabcentral/fileexchange/31926-fingerprint-minutiae-extraction?s_tid=srchtitle
- [62] M. Guri, "Air-viber: Exfiltrating data from air-gapped computers via covert surface vibrations," *arXiv preprint arXiv:2004.06195*, 2020.
- [63] *Long-range mmWave radar sensing demo*, <https://training.ti.com/mmwave-automotive-imaging-radar-system-long-range-detection>, Accessed: 2021-1-4.
- [64] "Sha-2," April 2021. [Online]. Available: <https://en.wikipedia.org/wiki/SHA-2>
- [65] M. Guri, Y. Solewicz, A. Daidakulov, and Y. Elovici, "Diskfiltration: Data exfiltration from speakerless air-gapped computers via covert hard drive noise," *arXiv preprint arXiv:1608.03431*, 2016.
- [66] R. Nandakumar, A. Takakuwa, T. Kohno, and S. Gollakota, "Covert-band: Activity information leakage using music," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, no. 3, pp. 1–24, 2017.
- [67] 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*, M. Polychronakis and M. Meier, Eds. Cham: Springer International Publishing, 2017, pp. 161–184.
- [68] M. Guri, O. Hasson, G. Kedma, and Y. Elovici, "Visisploit: An optical covert-channel to leak data through an air-gap," *arXiv preprint arXiv:1607.03946*, 2016.
- [69] M. Guri, G. Kedma, A. Kachlon, and Y. Elovici, "Airhopper: Bridging the air-gap between isolated networks and mobile phones using radio frequencies," in *2014 9th International Conference on Malicious and Unwanted Software: The Americas (MALWARE)*. IEEE, 2014, pp. 58–67.
- [70] B. Zhao, M. Ni, and P. Fan, "Powermitter: data exfiltration from air-gapped computer through switching power supply," *China Communications*, vol. 15, no. 2, pp. 170–189, 2018.
- [71] 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*, 2016.
- [72] R. Zhao, T. Woodford, T. Wei, K. Qian, and X. Zhang, "M-cube: a millimeter-wave massive mimo software radio," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020, pp. 1–14.
- [73] S. Jog, J. Wang, J. Guan, T. Moon, H. Hassanieh, and R. R. Choudhury, "Many-to-many beam alignment in millimeter wave networks," in *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*, 2019, pp. 783–800.
- [74] Y. Zhu, Y. Zhu, Z. Zhang, B. Y. Zhao, and H. Zheng, "60ghz mobile imaging radar," in *Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications*. ACM, 2015, pp. 75–80.
- [75] Y. Zhu, Y. Yao, B. Y. Zhao, and H. Zheng, "Object recognition and navigation using a single networking device," in *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 2017, pp. 265–277.
- [76] C. X. Lu, S. Rosa, P. Zhao, B. Wang, C. Chen, J. A. Stankovic, N. Trigoni, and A. Markham, "See through smoke: robust indoor mapping with low-cost mmwave radar," in *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 2020, pp. 14–27.
- [77] C. X. Lu, M. R. U. Saputra, P. Zhao, Y. Almalioglu, P. P. de Gusmao, C. Chen, K. Sun, N. Trigoni, and A. Markham, "milliego: single-chip mmwave radar aided egomotion estimation via deep sensor fusion," in *Proceedings of the 18th Conference on Embedded Networked Sensor Systems (SenSys)*, 2020.
- [78] J. Guan, S. Madani, S. Jog, S. Gupta, and H. Hassanieh, "Through fog high-resolution imaging using millimeter wave radar," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [79] Y. Zhu, Y. Zhu, B. Y. Zhao, and H. Zheng, "Reusing 60ghz radios for mobile radar imaging," in *MobiCom*, 2015.
- [80] B. Chen, H. Li, Z. Li, X. Chen, C. Xu, and W. Xu, "Thermowave: A new paradigm of wireless passive temperature monitoring via mmwave sensing," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. ACM, 2020.
- [81] Z. Li, Z. Yang, C. Song, C. Li, Z. Peng, and W. Xu, "E-eye: Hidden electronics recognition through mmwave nonlinear effects," in *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*, 2018, pp. 68–81.
- [82] A. Prabhakara, V. Singh, S. Kumar, and A. Rowe, "Osprey: a mmwave approach to tire wear sensing," in *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 2020.
- [83] Z. Li, F. Ma, A. S. Rathore, Z. Yang, B. Chen, L. Su, and W. Xu, "Wavespy: Remote and through-wall screen attack via mmwave sensing," in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 217–232.
- [84] Z. Li, B. Chen, Z. Yang, H. Li, C. Xu, X. Chen, K. Wang, and W. Xu, "Ferrotag: A paper-based mmwave-scannable tagging infrastructure," in *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*, 2019, pp. 324–337.
- [85] F. Lin, C. Song, Y. Zhuang, W. Xu, C. Li, and K. Ren, "Cardiac scan: A non-contact and continuous heart-based user authentication system," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*. ACM, 2017, pp. 315–328.
- [86] C. Xu, Z. Li, H. Zhang, A. S. Rathore, H. Li, C. Song, K. Wang, and W. Xu, "Waveear: Exploring a mmwave-based noise-resistant speech sensing for voice-user interface," in *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*, 2019, pp. 14–26.
- [87] C. Jiang, J. Guo, Y. He, M. Jin, S. Li, and Y. Liu, "mmvib: micrometer-level vibration measurement with mmwave radar," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020, pp. 1–13.
- [88] H. Li, C. Xu, A. S. Rathore, Z. Li, H. Zhang, C. Song, K. Wang, L. Su, F. Lin, K. Ren et al., "Vocalprint: exploring a resilient and secure voice authentication via mmwave biometric interrogation," in *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, 2020, pp. 312–325.
- [89] Y. Zheng, Y. Zhang, K. Qian, G. Zhang, Y. Liu, C. Wu, and Z. Yang, "Zero-effort cross-domain gesture recognition with wi-fi," in *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*, 2019, pp. 313–325.
- [90] T. Wei and X. Zhang, "mtrack: High-precision passive tracking using millimeter wave radios," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, 2015, pp. 117–129.