

DiskSpy: Exploring a Long-Range Covert-Channel Attack via mmWave Sensing of μm -level HDD Vibrations

Weiye Xu^{†‡}, Danli Wen[†], Jianwei Liu^{†¶}, Zixin Lin[†], Yuanqing Zheng[§], Xian Xu[†], Jinsong Han^{†*}
[†]Zhejiang University, [§]The Hong Kong Polytechnic University,
[‡]China Mobile Research Institute, [¶]Hangzhou City University

{xuweiye, wendanli, jianweiliu, xianxu, hanjinsong}@zju.edu.cn, csyqzheng@comp.polyu.edu.hk

Abstract

An air-gapped environment is widely regarded as a secure measure against the leakage of sensitive information, as it is physically isolated from insecure external networks. This paper presents a new covert-channel attack named DiskSpy, which reveals the risk of secretly sending sensitive information from air-gapped environments by modulating hard disk vibrations. In particular, DiskSpy leverages the vibrations of commonly used storage devices, hard disk drives (HDDs), in air-gapped computers to encode sensitive information. It then employs millimeter-wave (mmWave) to sense these vibrations and decode the underlying data. In practice, HDD vibrations are extremely weak and mmWave signals suffer significant power attenuation in long-distance propagation. To realize a practical attack at a long distance, we develop a novel mmWave-based long-range μm -level vibration sensing technique to push the limit of mmWave sensing. We implement DiskSpy with commercial off-the-shelf (COTS) mmWave radars and conduct extensive experiments. The experimental results show that even at a long attack range of 22m, DiskSpy can send secret information to a remote mmWave radar at 20bps with a BER lower than 1.2%. More importantly, DiskSpy has no restriction on the mounting manner and placement of the HDD, and can launch attacks even in the non-line-of-sight (NLOS) scenarios.

1 Introduction

With the prevalence of network threat techniques (e.g., ransomware [36], phishing [25], DDoS attacks [46] and advanced persistent threats [7]), security-conscious organizations employ various measures to prevent the exfiltration of sensitive information. Among these measures, establishing an air-gapped network is believed highly effective and widely adopted in many critical infrastructure sectors [33] such as governmental computer systems and financial computer networks. In air-gapped environments, computers storing sensi-

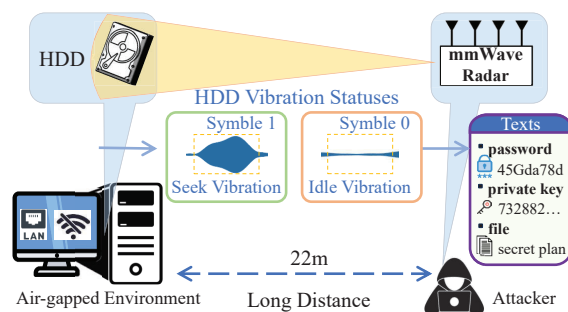


Figure 1: DiskSpy exfiltrates sensitive data from an air-gapped computer by modulating its HDD vibrations and sensing the vibrations with a mmWave radar.

tive information are physically isolated from unsecured networks and unauthorized users [12, 31].

Although air-gapped networks can enhance security, they are not immune to information leakage. For example, previous studies have demonstrated that air-gapped computers could be breached by covert-channel attacks. These attacks inject malicious software into air-gapped computers and secretly leak sensitive information to attackers (e.g., acoustic emission from fan [19], RF leakage from display cable [16], thermal emanation [4, 17]). The covertness of such channels makes it difficult for security systems to detect and defend.

Existing covert-channel attacks targeting air-gapped computers typically have the following drawbacks that limit their impact on practical air-gapped systems. First, to build a covert channel, attackers often exhibit abnormal behaviors (e.g., a full load of CPU [17, 21] or memory [15]), which raise suspicion of victims and can be easily detected. Second, due to the properties of the media involved in constructing covert channels, some covert channels (e.g., optical signal [14, 20]) require line-of-sight (LOS) paths to compromised computers [17, 19, 31]. Furthermore, owing to the signal attenuation, current covert-channel attacks [13, 17, 18] typically require proximity to compromised devices, such as contact-based or

*Jinsong Han is the corresponding author.

short-range sensing. Nevertheless, as the security-essential organizations usually establish a strict boundary around confidential areas, strictly prohibiting unauthorized devices [43], these short-range covert-channel attacks can be effectively detected and mitigated. These shortcomings limit their impact on real-world air-gapped security systems.

By developing DiskSpy, this paper presents a new type of covert-channel attack that could potentially pose a serious security threat to existing air-gapped systems. As illustrated in Fig. 1, DiskSpy secretly leaks the sensitive information from an air-gapped computer by modulating the HDD vibration of a compromised computer, while an attacker decodes the leaked information by sensing and demodulating the vibration with a COTS mmWave radar. As HDDs are widely used for mass storage in air-gapped computers and mmWave radars are readily available on the market, such a covert channel could have a profound impact on security systems in operation.

We found building a covert channel using COTS HDD and mmWave radar is indeed possible yet challenging. In particular, such a covert channel poses an extremely high-performance requirement in terms of sensing resolution (e.g., $< 10\mu\text{m}$ vibration) and sensing range (e.g., $> 20\text{m}$). Nevertheless, current state-of-the-art mmWave-based vibration sensing works [5, 8, 48, 52] primarily focus on improving sensing precision over a limited sensing range (1-5m). Obviously, DiskSpy’s sensing requirements are beyond the sensing capability of existing solutions. *The required sensing capability is comparable to detecting the tiny motion of an ant (μm -level) on the ground from the top of an 8-story building ($> 20\text{m}$).* DiskSpy develops a series of novel techniques to covertly send secret information by modulating HDD vibrations, and sense and decode such information with mmWave sensing.

(1) How to modulate the HDD vibrations to establish a stealthy and efficient covert channel. The covert channel encodes sensitive information by modulating the vibration statuses of the HDD. Ideally, the modulated vibrations should assemble regular HDD vibrations for stealthiness and can be rapidly changed for modulation efficiency. Based on the characterization of HDDs, DiskSpy controls HDDs to generate different vibration patterns with the seek operations of HDD to encode sensitive information. Since the seek operations are routine HDD operations and can be invoked frequently with short intervals, DiskSpy can achieve a sufficiently high bit rate for information leakage.

(2) How to sense the subtle HDD vibrations over long distances and in NLOS scenarios using mmWave signals. mmWave signals suffer significant free-space path loss, which poses challenges in providing sufficient strength to sense the vibration over long distances. Moreover, in NLOS scenarios, obstacles like furniture and walls can block and reflect mmWave signals leading to signal attenuation by several orders of magnitude greater than those in LOS scenarios. To solve this problem, we develop two strategies that work jointly at both the transmitter (Tx) and receiver (Rx) ends of the

mmWave radar to meet the performance requirement of long-range and NLOS vibration sensing. i) At the Tx of radar, we propose an improved Tx beamforming method which can effectively enhance the SNR of vibration signals, thus improving the sensing range. Different from the traditional communication-oriented Tx beamforming method [11], our method is oriented by sensing. It can steer the formed beam to a direction that is optimal for vibration sensing rather than for communication. Meanwhile, since the formed beam concentrates energy and is directional, the beam can be steered and reflected by objects in the environment to achieve NLOS sensing of the target HDD. ii) At the Rx of radar, we propose the phase coherent integration of multiple Rx antennas to further amplify vibration signals caused by HDD modulations. The two strategies working together empower DiskSpy with unprecedented sensing resolution and range.

We implement DiskSpy solely with COTS components such as TI mmWave radars and comprehensively evaluate its performance under various experiment settings. The results with various types of HDDs from four mainstream manufacturers demonstrate that DiskSpy can achieve a sufficiently high data rate of 20bps with a Bit Error Rate (BER) below 1.2%. Furthermore, DiskSpy adapts to various HDD mounting and placement (e.g., inside and outside of the chassis). Moreover, DiskSpy is applicable to NLOS attack scenarios over a significantly increased distance of up to 22m.

Our contributions are summarized as follows.

- We reveal the risk of a first-of-its-kind covert channel, which leverages mmWave signals to sense modulated HDD vibrations for leaking information from air-gapped computers.
- We develop a Tx-Rx co-optimization method to push the limit of mmWave sensing in terms of sensing precision and range. To our knowledge, we are the first to achieve μm -level vibration sensing at a long range of 22m with COTS mmWave radar.
- We implement a proof-of-concept system with COTS HDDs and mmWave radars, and conduct extensive experiments in various settings. Results demonstrate that DiskSpy can leak secret data from air-gapped computers, which poses serious threats to security systems.

2 Threat Model

Attack model. We consider an attack model as follows: Alice has a computer containing sensitive information in an air-gapped network. The attacker, Bob, attempts to steal sensitive data from Alice’s computer outside the air-gapped environment. To this end, Bob needs to stay far away ($>20\text{m}$) from Alice and perform data acquisition without raising Alice’s awareness. Bob uses a mmWave radar to sense the vibrations

of the HDD on Alice’s air-gapped computer and extract the modulated information.

Practicality consideration. Given the practicality of the attack scenario, Bob is subject to the following constraints. (1) No proximity: Alice is vigilant against shoulder-surfing [42] or other attacks that can be observed through visual inspection. Therefore, Bob needs to stay as far away from the air-gapped computer as possible, let alone physically touch it or modify its hardware. (2) No physical connection: Alice disconnects the protected computer from external networks.

Attacker capability. Similar to the assumptions of conventional covert-channel attacks [21], we assume that Bob has the ability to install malware (only running in user space without requiring root privilege) with DiskSpy capability on the target computer. The malware is set to trigger at a specific time each day to launch the covert-channel attack. This can be accomplished via many attack vectors, such as the time-based Logic Bombs [6, 30]. Furthermore, we also assume that the malware can get access to certain sensitive data in the air-gapped network. This can be achieved by using various methods such as malicious insiders, supply chain attacks, and exploiting side-channel [32] [1]. The data will then be modulated by HDD vibrations and sensed by a mmWave radar. Moreover, we assume Bob has no prior knowledge about the targeted HDD, regardless of type, mounting manners, or vibration frequency.

Attack scenarios. Compared with existing covert-channel attacks, DiskSpy can support the following three challenging attack scenarios. (1) **Long-range.** Security-sensitive organizations usually establish a boundary to protect computers storing sensitive information, strictly prohibiting unauthorized devices. This boundary successfully defends against conventional short-range attack methods but falls short in countering DiskSpy. DiskSpy, with long-range attack capability, can conduct data infiltration outside the boundary. (2) **Non-line-of-sight.** Different from traditional visible light, DiskSpy makes new definitions of LOS and NLOS for mmWave sensing. Specifically, LOS is the scenario in which mmWave signals can directly reach the target without reflection, while NLOS is the scenario in which the LOS path is obstructed by some materials impenetrable to mmWave, such as metal. It is worth noting that when the obstacles are penetrable by mmWave signals, such as thin wood and glass, these scenarios are still categorized as LOS. Air-gapped computers are typically located in complex environments where LOS paths may be completely blocked. In these intricate NLOS scenarios, DiskSpy can still effectively conduct a covert-channel attack. (3) **Unknown air-gapped device’s location.** In practice, due to the random deployment and mobility of devices, the exact position of the air-gapped device may be unknown to the attacker, thereby escalating the challenge of launching an attack. In this case, DiskSpy can still exfiltrate data even when the HDD’s location is unknown. This is achieved by exploiting the SNR of the HDD vibration signal to guide the mmWave beam towards the HDD via either LOS or reflection paths.

3 Background

3.1 Hard Disk Driver

HDDs are widely used for mass storage in computers. Fig. 2 illustrates the structure of a typical HDD. A platter, coated with magnetic material, stores data and rotates at a certain speed. Data is read from or written to the platter by detecting or changing the magnetization on its surface through the R/W head. The surface of the platter is composed of circular tracks that store data. Each track consists of multiple sectors, which are the smallest storage units on the HDD. During seek operations, the HDD moves the actuator arm and the R/W head in an arc to locate the target track. These movements cause HDD vibration. Previous works [47] utilize the capacitive probe to measure the amplitudes of HDD vibrations, revealing they are typically less than $10\mu\text{m}$. In DiskSpy, we exploit these subtle vibrations to encode secret information.

3.2 Vibration Sensing with mmWave

DiskSpy employs mmWave to capture the HDD vibration. Generally, a mmWave radar transmits a chirp signal, whose frequency increases linearly with time. The chirp signal can be formulated as $S_{T_x}(t) = e^{j2\pi(f_c t + \frac{Kt^2}{2})}$, where K and f_c are the slope of frequency and the starting frequency, respectively. Assuming that there is a vibrating target (e.g., an HDD) and the time-variant distance between the HDD and the mmWave radar is denoted as $R(t)$, the received signal can be expressed as: $S_{R_x}(t) = \alpha e^{j2\pi(f_c(t - (2R(t)/c)) + \frac{K(t - (2R(t)/c))^2}{2})}$, where c and α are the speed of light and the path loss, respectively.

The received signal will be mixed with the transmitted signal. Then, we can get an intermediate frequency (IF) signal [34]: $S(t) = \alpha S_{T_x}^*(t) S_{R_x}(t) \approx \alpha e^{j(\frac{4\pi K R(t)}{c}t + \frac{4\pi f_c R(t)}{c})}$. Next, we perform fast Fourier transform (Range-FFT) [39] on the IF signal $S(t)$ to measure the signal $s_r(t)$ reflected by the vibrating target:

$$s_r(t) = \alpha \exp(j(4\pi f_c R(t)/c)). \quad (1)$$

By tracking the phase changes, the mmWave radar can measure the vibration of the target object.

4 Characterizing Covert Channel

An HDD exhibits various vibration statuses under different working modes. In this section, we first analyze the vibration statuses of HDD. Then, we test the feasibility of sensing such vibrations with mmWave to construct a covert channel.

4.1 Vibration Statuses of HDD

The vibration statuses of an HDD can be classified into the following categories [29] [47]: static, idle vibration, and seek




Figure 2: Structure of a hard disk driver.




Figure 3: Vibration frequency of the static HDD.




Figure 4: Idle vibration frequency distribution.




Figure 5: Seek vibration frequency distribution.

vibration. HDD is **Static** when the HDD is inactive (e.g., sleep mode). In this status, both the HDD platter and R/W head remain stationary. As shown in Fig. 3, the mmWave signals reflected by a static HDD show no vibration. **Idle vibration** occurs when the HDD platter spins while the R/W head remains stationary. Due to the hardware imperfection (e.g., imbalance of the rotating platter), an HDD exhibits idle vibration at a frequency corresponding to the rotational speed of the platter [29], which can be calculated by $RPM/60$ (RPM is the revolutions per minute of the HDD). As shown in Fig. 4, we utilize a mmWave radar to capture the idle vibration generated by an HDD spinning at 7200 RPM. With the frequency domain analysis on the phase of mmWave signals, we can detect an idle vibration at the frequency of $f_0 = 7200/60 = 120Hz$. It is worth noting that the RPM of an HDD (corresponding to the frequency of idle vibration) is typically stable when the HDD is in active mode. **Seek vibration** can be generated by the seek operation of the HDD. When the actuator moves the R/W head to a specific track to access data, the acceleration of the head and the applied force cause a slight vibration of the HDD [29]. As shown in Fig. 5, we set the HDD to perform seek operation repeatedly over a period, e.g., oscillating the R/W head between sector A and sector B, and utilize the mmWave radar to sense the seek vibration. We find that along with the idle vibration at the frequency of $f_0 = 120Hz$, the seek operations also generate an additional frequency component f_1 .

A simple yet inefficient solution to build a covert channel can be alternating between the static and the idle vibration. However, this strawman solution has two key limitations: (1) Low efficiency. It requires the HDD to alternate between inactive and active mode. Waking up the HDD from inactive mode usually takes a long time (e.g., a few seconds), which limits the data rate. (2) Lack of stealthiness. Setting the HDD to inactive mode will interrupt the normal use of the HDD. Instead, we build a covert channel by alternating between the idle vibration and the seek vibration.

4.2 Sensing HDD Vibrations with mmWave

Harnessing idle and seek vibrations for a covert-channel attack requires accurate detection of these vibrations with

mmWave. To test its feasibility, we start with a simplified experiment and then elaborate on our design considerations.

4.2.1 Modeling Vibration Sensing

In the HDD-mmWave covert channel, an HDD is positioned in front of the mmWave radar at an angle of θ_0 and the distance between the HDD and radar is R_0 , as shown in Fig. 6. When the HDD generates seek vibrations, according to Eq. 1, the phase of the received IF signal reflected from the HDD can be represented as:

$$\phi(t) = \left(\frac{4\pi f_c}{c} (R_0 + \underbrace{v(t)}_{v'(t)} \cdot \cos\alpha) + \phi_{noise} \right) \bmod 2\pi. \quad (2)$$

ϕ_{noise} is the noise in the phase value. $v(t)$ is the time-varying micro-displacement produced by the seek vibration. α is the misaligned angle between the seek vibration and the radar sensing directions. This means the vibration displacement $v'(t)$ measured by the mmWave radar is the projection of $v(t)$ along the sensing direction. According to Sec. 4.1, the seek vibration has two primary frequency components, so $v(t)$ can be represented as:

$$v(t) = A_0 \sin(2\pi f_0 t) + A_1 \sin(2\pi f_1 t), \quad (3)$$

where A_0 and A_1 represent the amplitudes of vibrations caused by platter rotation and R/W head movement, respectively. According to Eq. 2 and Eq. 3, the value of phase ϕ can reveal the vibration amplitude. However, since both A_0 and A_1 are tiny, i.e., smaller than $10\mu m$ [47], it is difficult to directly detect and differentiate idle and seek vibrations based on amplitudes. Fig. 7(a) plots the vibration signal of an HDD (7200 RPM) captured by the radar when the HDD executes the “seek-idle-seek” command. It is hard to distinguish these two vibration statuses by measuring the amplitudes.

4.2.2 Sensing HDD Vibrations Based on Frequency

We aim to detect and differentiate different vibration statuses through frequency analysis. According to Sec. 4.1, idle vibration exhibits a single primary frequency component f_0 , while seek vibration has two components, f_0 and f_1 . Consequently,




Figure 6: Model of the HDD vibration.




Figure 7: Raw and filtered vibrations.




Figure 8: f_1 exploration.

the distinction between these two vibrations hinges on the detection of f_1 . Unlike the fixed platter rotation frequency f_0 , the frequency f_1 aroused by R/W head movement is affected by many factors. In the following, we will explore f_1 by analyzing the time it takes to move the R/W head to the desired track (i.e., the seek time). The seek time T_s consists of two parts: the time T_b to start the actuator arm and the time to move the R/W head to the desired track. We can calculate the seek time as $T_s = T_b + m * T_t$, where m is the number of tracks to move and T_t is the time required to move the R/W head between adjacent tracks. Among these parameters, T_b and T_t are primarily determined by the disk type. m depends on the start and the end positions of the R/W head. Therefore, the seek time T_s is not constant but depends on various factors. Essentially, for different HDDs and different seek commands, the frequency f_1 can be different.

To address this issue, we explore an adaptive method to find the R/W head movement frequency f_1 . Specifically, assuming that we know the statuses of HDD vibration switching within a short period of time, e.g., “seek-idle-seek”. We collect the mmWave signal during this period and perform the Short-Time Fourier Transform (STFT) to generate the spectrogram, as shown in Fig. 8. This spectrogram can reveal the energy distribution of the vibration signals in different frequency bands. When the HDD performs the “seek-idle-seek” vibration sequence, the vibration energy corresponding to the R/W head movement frequency f_1 will also have a corresponding “high-low-high” sequence. Therefore, the key question becomes how to reliably detect a frequency band (with a varying f_1) whose time-variant energy resembles a predefined pattern (e.g., “seek-idle-seek”). As shown in Fig. 8, we set a small frequency window with an interval of 2Hz. Then, we slide the frequency window along the spectrogram. For each frequency window, at each time sampling point, we sum the energy associated with all frequencies within that window. The summed energy represents the energy level of the HDD vibration within that frequency window. We apply a typical matched filter [2] with the specified pattern (e.g., “high-low-high” sequence) to the energy level of each frequency window to infer the f_1 location. Fig. 8 shows the output after applying the matched filter, denoted as correlation value, which can reveal the correlation between the input energy level and the specified pattern. Thus, the frequency corresponding to the highest correlation value ρ_1 can be regarded as f_1 .

To demonstrate the effectiveness of our frequency determination method, we use it to determine f_1 of the aforementioned “seek-idle-seek” vibration sequence. Then, we exploit a bandpass filter with its lower and upper stopping frequencies set according to f_1 , to process the signals. Fig. 7(b) shows the results. Compared with the raw vibration signal (Fig. 7(a)), the vibration caused by the movement of the R/W head is more prominent. Meanwhile, we find that the seek and idle vibrations can be clearly distinguished by frequency analysis.

5 DiskSpy Design

5.1 DiskSpy Architecture

We propose DiskSpy, a long-range covert-channel attack that can exfiltrate sensitive data on an air-gapped computer from its HDD via mmWave sensing. As illustrated in Fig. 9, at a high level, the architecture of DiskSpy consists of two modules: the initiator (enabled by the HDD) and the responder (enabled by the mmWave radar).

At the initiator end, DiskSpy first gets access to the sensitive data on the air-gapped computer via infiltrated malicious code. Then, the data will be encoded into HDD vibrations (i.e., the transitions between the idle and seek vibration statuses) according to our information encoding scheme and transmitted in the form of a packet.

At the responder end, DiskSpy utilizes a mmWave radar to continuously interrogate the HDD to remotely sense its vibration statuses. Subsequently, the captured vibration statuses will be translated into sensitive information based on our proposed vibration decoding scheme.

However, conducting a long-range covert-channel attack via remote sensing of HDD vibrations is challenging due to the significant free-space path loss of the mmWave signal and the extremely low amplitude of HDD vibrations. To this end, we propose a long-range subtle vibration sensing approach (Sec. 5.2) and then design DiskSpy’s initiator and responder based on this approach (Sec. 5.3).

5.2 Long-range Subtle Vibration Sensing

The key to conducting a long-range covert-channel attack is to ensure a high SNR of the recovered vibration signals. To this end, DiskSpy proposes a joint optimization strategy that

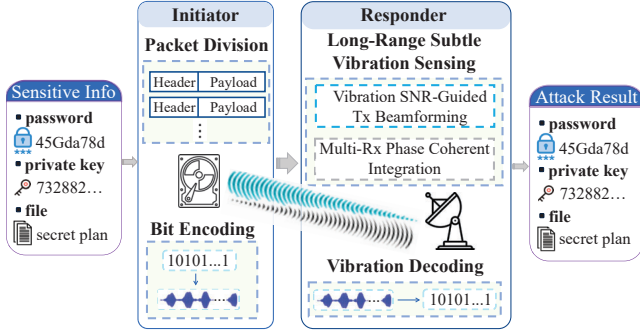


Figure 9: Architecture of DiskSpy.

improves the SNR of vibration signals at both the Tx and Rx ends of the mmWave radar.

5.2.1 Key Challenges

According to Eq. 2, the seek vibration $v'(t)$ can be measured by the phase (ϕ) of the received mmWave signal $S(t)$. To accurately sense the seek vibration, it is crucial to ensure that the SNR of the measured vibration signal $v'(t)$ (denoted as $SNR_{v'(t)}$) is sufficiently high. **However, as the attack distance increases, detecting seek vibrations becomes more challenging due to the decrease of $SNR_{v'(t)}$.** Specifically, $SNR_{v'(t)}$ can be calculated as the ratio of the power of the measured vibration signal to that of noise, i.e., $SNR_{v'(t)} = \frac{P(v'(t))}{P(noise)}$. Here, the noise power $P(noise) = P(N_{S(t)}) + P(N_{other})$ consists of two components. $N_{S(t)}$ denotes the noise arising from the channel noise of $S(t)$. N_{other} represents other noises unrelated to $S(t)$ channel, such as the noise caused by imperfect phase estimation owing to the limited resolution of Range-FFT [9, 28]. As the attack distance increases, $S(t)$ suffers free-space path loss [39], leading to an exponential growth of the channel noise of $S(t)$ and its corresponding noise power $P(N_{S(t)})$. In this case, the value of $SNR_{v'(t)}$ will decrease significantly, thereby posing a substantial challenge in conducting a long-range attack.

To further illustrate the impact of the attack distance on our covert channel, we compare the performance of recovering seek vibration under different attack distances. Fig. 10 shows the recovered HDD vibration when the attack distance is set to 1, 18, and 22m, respectively. The HDD performs the “seek-idle-seeking” vibration sequence. It can be observed that when the distance increases to 18m (Fig. 10(b)) or 22m (Fig. 10(c)), the measured phases of seek vibration are completely overwhelmed by the noise.

To address this challenge, we propose a long-range subtle vibration sensing technology encompassing two strategies for decreasing the aforementioned two noises $N_{S(t)}$ and N_{other} to enhance $SNR_{v'(t)}$. These two strategies are implemented at the Tx end and the Rx end of the mmWave radar, respectively.

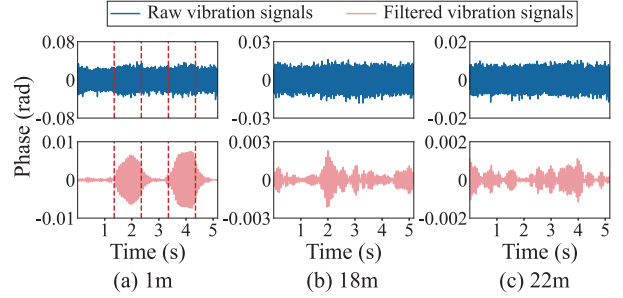


Figure 10: Vibration signals captured in normal mode without any vibration SNR improvement at different attack distances.

They work jointly to optimize the SNR of vibration signals, thus enabling long-range subtle vibration sensing.

5.2.2 Tx-end Optimization

At the Tx end of mmWave radar, we focus on mitigating the noise $N_{S(t)}$ to improve $SNR_{v'(t)}$. To this end, we first analyze the impact of $N_{S(t)}$ on $SNR_{v'(t)}$ and then describe the method to suppress this impact.

Impact of $N_{S(t)}$ on $SNR_{v'(t)}$. According to the typical parameter estimator Cramer Rao lower bounds [9], the power of $N_{S(t)}$ can be calculated as: $P(N_{S(t)}) = \frac{c^2}{16\eta\pi^2 f_c^2 SNR_{S(t)}}$, where η denotes the sampling coefficient and it is regarded as a constant in our work. $SNR_{S(t)}$ represents the SNR of the mmWave signal channel $S(t)$. It is worth noting that the difference between $SNR_{S(t)}$ and $SNR_{v'(t)}$ is that $SNR_{v'(t)}$ measures the SNR of the vibration signal $v'(t)$ extracted from $S(t)$. In other words, $SNR_{v'(t)}$ assesses the quality of the phase of $S(t)$. As the attack distance increases, $SNR_{S(t)}$ will decrease, thus increasing the value of the noise power $P(N_{S(t)})$. A high noise interference leads to a low $SNR_{v'(t)}$, making it difficult to accurately sense the subtle vibration signal. Therefore, to maintain a high $SNR_{v'(t)}$ under long-range attack conditions, it is essential to enhance $SNR_{S(t)}$.

Vibration SNR-guided Tx beamforming. An intuitive way to enhance $SNR_{S(t)}$ is to employ traditional Tx beamforming technology [38]. In the Tx beamforming mode, multiple antennas emit signals simultaneously to form a narrow high-energy beam. The formed beam is then steered towards an optimal direction to amplify the signal energy at the target location. To find the optimal steering direction, the standard beam steering method will sweep the beam in the space to find the direction with the strongest reflection, i.e., the one with the highest $SNR_{S(t)}$. However, this traditional method cannot be directly applied to enhance $SNR_{v'(t)}$, due to the misalignment of sensing direction (the direction of $S(t)$) and the vibration direction (the direction of $v(t)$).

To provide an in-depth explanation of this issue, we build a vibration sensing model by analyzing the signal propagation




Figure 11: Vibration sensing model.

process, as shown in Fig. 11. Without losing the effectiveness of modeling, taking the simplest LOS scenario as an example, the traditional Tx beamforming method will directly steer the beam towards the target object to obtain the best $SNR_{S(t)}$. However, in this case, there is an angle $\alpha = f(\theta_0)$ between the vibration direction and sensing direction (denoted as θ_0 in Fig. 11). As the Tx sweeps the beam to find the optimal beam direction, i.e., changes the angle θ_0 , α will vary accordingly, altering the value of measured vibration signal $v'(t) = v(t) \cdot \cos\alpha$. This implies that variations in θ_0 will simultaneously affect two factors of $SNR_{v'(t)}$, i.e., $P(v'(t))$ and $P(N_{S(t)})$. Traditional Tx beamforming method only finds the optimal direction θ_0 with the minimum value of $P(N_{S(t)})$ without considering the impact on $P(v'(t))$. Therefore, the direction identified by the traditional method may not correspond to the optimal $SNR_{v'(t)}$. Then, we conduct a verification experiment. As shown in Fig. 12, we steer the Tx beam along two paths towards the HDD, respectively. Path 1 follows the LOS path and Path 2 follows the reflection path. In Path 1, there is an angle α between the sensing direction and the vibration direction, while in Path 2, $\alpha = 0$. Fig. 12(b) shows $SNR_{S(t)}$ under these two different paths. It can be observed that, due to the smaller path loss in Path 1, the signal strength of $S(t)$ corresponding to Path 1 is higher than that of Path 2, i.e., $SNR_{S(t)}$ of Path 1 is higher than that of Path 2. Therefore, the traditional Tx beamforming method will choose Path 1 as the optimal direction. However, according to Fig. 12(c), the quality of the vibration signal recovered by the mmWave signal along Path 1 is lower than that along Path 2, i.e., $SNR_{v'(t)}$ of Path 1 is lower than that of Path 2. Therefore, traditional Tx beamforming is not applicable to our work.

To this end, we propose an improved Tx beamforming method guided by $SNR_{v'(t)}$. This method takes into account the impact of both $P(v'(t))$ and $P(N_{S(t)})$ on $SNR_{v'(t)}$ simultaneously. Specifically, when launching a long-range attack to an HDD with an unknown location in a new scenario, the attacker can sweep the direction of the Tx beam (θ) and calculate $SNR_{v'(t)}$ for each θ . Among all directions, the direction with the highest $SNR_{v'(t)}$ will be regarded as the optimal angle for sensing. The calculation of $SNR_{v'(t)}$ involves three steps. Firstly, we perform Range-FFT on the received mmWave signals and extract the phase $\phi_{\theta, R_i}(t)$ for each range bin R_i . Then, we extract the vibration signal $v'_{\theta, R_i}(t)$ from $\phi_{\theta, R_i}(t)$




Figure 12: The $SNR_{S(t)}$ along Path 1 is higher than along Path 2, while the $SNR_{v'(t)}$ along Path 1 is lower than along Path 2.

according to Eq. 2. Finally, according to the HDD vibration frequency, e.g., f_0 , we apply a bandpass filter (BPF) to $v'_{\theta, R_i}(t)$ and calculate its vibration SNR as

$$SNR_{v'_{\theta, R_i}(t)} = \frac{P(BPF(v'_{\theta, R_i}(t), f_0))}{P(v'_{\theta, R_i}(t)) - P(BPF(v'_{\theta, R_i}(t), f_0))}. \quad (4)$$

When applying the vibration SNR-guided Tx beamforming approach in the real attack scenario, we choose the idle vibration of the HDD as the one measured by $SNR_{v'(t)}$. This is because the idle vibration always exists, requiring no manual intervention. In this case, the cutoff frequencies of the bandpass filter are set according to the typical RPMs of modern consumer-grade HDDs.

With the help of vibration SNR-guided Tx beamforming, the attack distance can be significantly extended. As shown in Fig. 13, the vibration statuses of the HDD can be accurately captured even when the attack distance is 18m.

5.2.3 Rx-end Optimization

To detect the tiny vibration at a long-range, merely decreasing the impact of noise from $N_{S(t)}$ may not be sufficient as illustrated in Fig. 13(c). To further improve the sensing distance, we also need to mitigate the impact of another noise component N_{other} , as illustrated in Sec. 5.2.1. The noise N_{other} is generated during the phase estimation process. Specifically, according to Sec. 3.2, we perform Range-FFT on the received mmWave signal to estimate the phase at the target bin for vibration signal recovery. Nevertheless, due to the limited resolution of the Range-FFT, the residual phase caused by imperfect target bin estimation introduces noise to the measured vibration signal, i.e., N_{other} . Concerning the subtle vibrations of HDD, the minor impact of the noise N_{other} may become substantial. Therefore, at the Rx end of radar, we propose the multi-RX phase coherent integration algorithm, which combines multiple Rx antennas to suppress the noise N_{other} .

The multi-RX phase coherent integration algorithm is designed based on our key observation that the vibration signals sensed by different Rx antennas exhibit coherence, i.e., they share the same initial phase. As shown in Fig. 11, due to the arrangement of Rx antennas, the mmWave signal reflected by the HDD and received

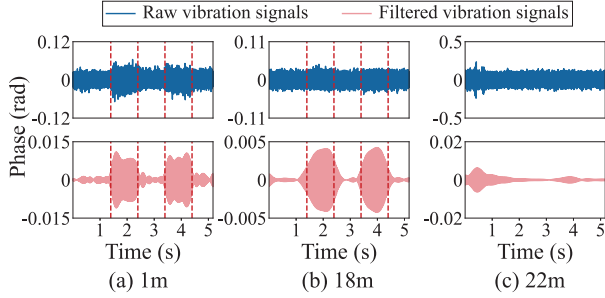


Figure 13: Vibration signals reinforced only by employing Tx optimization under different attack distances.

by Rx2 will travel an additional distance $d \sin \theta_0$ than the one received by Rx1. This additional distance leads to discrepancies in the sampling delay of the vibration signal, meaning that Rx1 and Rx2 sample the same vibration signal at different time instances. Therefore, according to Eq. 3, the HDD vibration recovered by the j -th Rx antenna can be represented as $v'_{\theta_0,j}(t) = A_0 \sin(2\pi f_0 t + \phi_{f_0,j}) + A_1 \sin(2\pi f_1 t + \phi_{f_1,j})$. $\phi_{f_0,j} = 2\pi f_0 (j-1) \cdot \frac{d \sin \theta_0}{c}$ and $\phi_{f_1,j} = 2\pi f_1 (j-1) \cdot \frac{d \sin \theta_0}{c}$ are the phase differences between the j -th Rx and the 1st Rx due to the sampling delay. Note that the vibration frequencies (both f_0 and f_1) of the HDD are around 100Hz, and the speed of light $c = 3 \times 10^8$ m/s. Therefore, the values of $\phi_{f_0,j}$ and $\phi_{f_1,j}$ are approximately 0, i.e., the initial phases of the vibration signals recovered by different Rx antennas are synchronized. To improve the $SNR_{v'(t)}$, inspired by the coherent integration [40], we accumulate the phase-synchronous vibration signals recovered by multiple Rx antennas. In this case, the coherent vibration signal power will be amplified by N^2 , while the random and irregular noise power increased by N . Therefore, after the coherent integration, $SNR_{v'(t)}$ can be calculated as $SNR_{v'(t)} = \frac{N^2 P(v'(t))}{NP(N_{other})} = N \frac{P(v'(t))}{P(N_{other})}$, i.e., $SNR_{v'(t)}$ has been amplified by N times. Different from conventional coherent integration in radar signal processing [51], we are the first to apply the coherent integration to the phase of mmWave signal rather than the mmWave signal itself, thus directly boosting the $SNR_{v'(t)}$. Furthermore, the conventional method often requires multiple signal collections within a specific time interval to obtain coherent signals. This implies a trade-off, sacrificing time to enhance the SNR. However, our research reveals the coherence in vibration signals obtained by different Rx antennas. This discovery enables us to achieve SNR improvement within a shorter time.

Fig. 14 depicts the effectiveness of recovering the HDD vibration statuses after employing vibration signal SNR improvement at both the Tx and Rx ends. Notably, DiskSpy exhibits excellent accuracy in capturing the tiny vibration of the HDD even when the attack distance is as long as 22m. This capability effectively supports the feasibility of practical

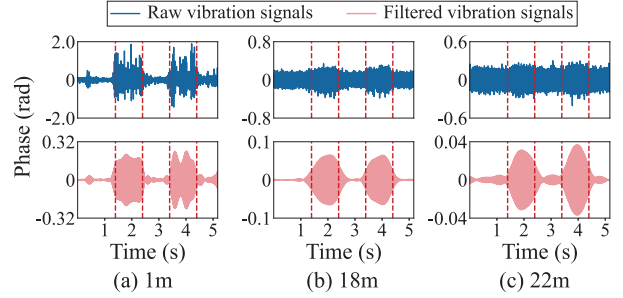


Figure 14: Vibration signals reinforced by employing Tx-Rx optimization under different attack distances.

long-distance covert-channel attacks.

5.3 Design of Initiator and Responder

5.3.1 DiskSpy Initiator

The goal of DiskSpy's initiator is to encode the sensitive data into HDD vibrations covertly. However, to establish a practical covert channel, DiskSpy's initiator needs to address two technical challenges. (1) Enabling vibration decoding. Since the seek vibration parameters (e.g., frequency and amplitude) are affected by the seek operations used in DiskSpy's initiator (according to Sec. 4.2.2), the vibration decoding parameters on DiskSpy's responder need to synchronize with the initiator dynamically. (2) Handling large-scale data. Extended transmission times due to large data size may increase the risk of system interruptions, thus compromising the transmission process. To solve these challenges, we propose a data encoding scheme, and its pseudocode is presented in Appendix A. This scheme designs a packet structure and divides the sensitive data into multiple packets for bit encoding.

Packet division. To facilitate vibration decoding, the packet structure consists of a packet header and a payload part. The header is set to "101" sequence, which serves three purposes. (1) Marking the beginning of the transmission. (2) Determining the frequency f_1 of R/W head movement. As mentioned in Sec. 4.2.2, the value of f_1 is affected by the seek operations utilized in bit encoding. The specific vibration pattern of the header can be utilized by DiskSpy's responder to determine f_1 . (3) Providing criteria for decoding. Based on the header, we can also calculate the vibration amplitudes of '1' and '0', which provide criteria for the subsequent data decoding. We will detail how to determine f_1 and the vibration amplitude in Sec. 5.3.2. The payload carries the sensitive data that needs to be transmitted and its length is fixed, denoted as $length_p$. If the length of the remaining untransmitted data is less than $length_p$, we will append zeros at the end of the payload.

To support large-scale data transmission, the data encoding scheme divides sensitive data into multiple packets for transmission and maintains a record of the current packet




Figure 15: Radar implementation.




Figure 16: Experiment setup.




Figure 17: HDD mounting manners.

being transmitted, i.e., $index_p$. In the event of a system interruption, the data transmission task will pause and preserve $index_p$. Upon task recovery, the transmission resumes from the recorded $index_p$, ensuring seamless continuity in the data transmission process. At the end of the data encoding, DiskSpy’s initiator will send a flag ($flag_{end}$) to inform DiskSpy’s responder that the data encoding process is completed and convey the total length of sensitive data.

Bit encoding. After packet division, we will encode bits in the packets into HDD vibrations. As introduced in Sec. 4.1, DiskSpy selects two natural HDD vibration statuses: idle vibration and seek vibration, for bit encoding. Specifically, we design the *bit encoding function*, i.e., $bitEncoding$, to seamlessly transition between these two statuses, facilitating effective data encoding. In $bitEncoding$, we employ on-off keying (OOK) modulation: the presence of the vibration aroused by R/W head movement (seek vibration) for a duration of T_1 represents bit 1, while its absence (idle vibration) for a duration of T_0 represents bit 0. Bit ‘1’/‘0’ can be generated by moving/fixing the R/W head with/without seek operation. It is worth noting that to enhance the stealthiness of DiskSpy, $bitEncoding$ uses files’ positions on the HDD as the source address ($address_s$) and destination address ($address_d$) for seek operations. Since these addresses are file positions rather than sector numbers, root permissions are not required. The implementation of the $bitEncoding$ function is described in Appendix A. When transmitting ‘1’, $bit_encoding$ invokes seek operations, repeatedly moving the R/W head between $address_s$ and $address_d$. To transmit ‘0’, the R/W head remains stationary. To ensure precise manipulation of the HDD statuses, we disable the disk cache, eliminating timing delays and inconsistencies in vibration generation.

5.3.2 DiskSpy Responder

At the responder end, DiskSpy first employs the long-range subtle vibration sensing approach to remotely sense HDD vibrations. Then, we design a vibration decoding method to decode the HDD vibration signals (formed in the packet structure) to get meaningful sensitive data. The vibration decoding method involves two steps: (1) detecting the packet header to determine the frequency f_1 and the amplitude of

the seek vibration, (2) decoding the payload section of the packet according to the obtained f_1 and the amplitude. Then, we present the details of the vibration decoding method.

To locate the beginning of data transmission, we first need to detect the packet header. As aforementioned, the vibration pattern of the header is “101”. Therefore, we set a time window W_{header} based on the duration of the header and slide it along the vibration signal to find the “101” vibration sequence. Specifically, in each W_{header} , we utilize the adaptive R/W head movement frequency determination method proposed in Sec. 4.2.2 to find the candidate f_1 and its correlation value ρ_1 . As a result, we obtain a candidate f_1 list and its correlation value list $[\rho_{1,1}, \dots, \rho_{1,k}, \dots, \rho_{1,N_{step}}]$, where $\rho_{1,k}$ represents the correlation value of the k -th W_{header} and N_{step} is the total sliding steps of W_{header} . Next, we find out the maximum correlation value ρ_{max} in the correlation value list and compare it with a pre-defined threshold ρ_{thre} . If ρ_{max} surpasses the threshold, the header is regarded as the W_{header} where we get ρ_{max} , and the corresponding candidate f_1 is the authentic R/W head movement frequency.

After getting the header and f_1 , we also need to determine the vibration amplitude of a ‘1’ and a ‘0’. Based on f_1 , we exploit a bandpass filter to extract the vibration aroused by R/W head movement. Since we adopt the OOK modulation, the amplitude of the vibration signal serves as a direct decoding indicator. Specifically, based on the amplitude of bit ‘1’ (A_{bit1}) and bit ‘0’ (A_{bit0}), we set an amplitude threshold $A_{thre} = (A_{bit1} + A_{bit0})/2$. If the majority (60%) of the vibration amplitude in a symbol duration is higher than A_{thre} , this symbol is decoded as ‘1’, otherwise, it is ‘0’.

6 Implementation

As shown in Fig. 15, we prototype DiskSpy with COTS mmWave radar components, including a mmWave radar board and a data acquisition board.

mmWave radar board. The mmWave radar board consists of four Texas Instruments (TI) AWR2243P [26] radar chips. Each chip contains three transmitting antennas and four receiving antennas. As shown in Fig. 15, all antennas on the radar board form two linear arrays: a Tx antenna array con-




Figure 18: Performance on various HDD models.

sisting of 12 uniformly spaced transmitting antennas and an Rx antenna array consisting of 16 receiving antennas. Tx antennas send 10000 FMCW chirp signals per second. For each chirp signal, the starting frequency is $f_c = 77\text{GHz}$ and the frequency slope is $K = 78.986\text{MHz}/\mu\text{s}$. To realize the vibration SNR-guided Tx beamforming, we first sweep the Tx beam between $[-90^\circ, 90^\circ]$ with a step of 2° to identify the optimal beam direction. Note that the beam sweeping only needs to be performed once for each attack scenario to detect the target HDD (the HDD location is unknown in advance), which takes around 18s in total. Then, utilizing the mmWave Studio, a software tool provided by TI, we can assign the initial phase for each Tx antenna and let all Tx antennas emit signals simultaneously, enabling the formation of a high-gain beam towards the HDD.

Data acquisition board. We capture the raw mmWave signals at high speed using the TI MMWCAS-DSP-EVM [27] data acquisition board. The captured data are then transmitted to a PC with Intel(R) Core(TM) i5-8250U CPU and 8 GB RAM. The data processing algorithms are implemented by Python and MATLAB.

7 Evaluation

7.1 Experimental Setting

Setup and data collection. As shown in Fig. 16, we evaluate the performance of DiskSpy in two typical usage scenarios of HDD: external connection to a computer and internal mounting within the chassis. Under these two scenarios, we also evaluate DiskSpy in four different HDD mounting manners (one in the external connection scenario and three in the internal mounting scenario as shown in Fig. 17). The target HDDs are located in a varied range up to 26m (HDD distance) away from the radar. The relative angle (HDD angle) between the HDD and the radar is denoted as θ varied in the range of $[-60^\circ, 60^\circ]$. At the initiator end of DiskSpy, we manipulate the vibration statuses of the HDDs to transmit secret data. By default, the HDD distance and angle are set to 10m and 0° , respectively. The time slots for encoding 0s and 1s are set




Figure 19: Performance on different attack distances and attack orientations.

to $T_0 = 25\text{ms}$ and $T_1 = 100\text{ms}$, respectively, with a bit rate of 16bps. We test DiskSpy with 9 different models of HDDs (listed in Fig. 18) from four mainstream manufacturers (Seagate, Western Digital, Toshiba and Hitachi). In each trial, we randomly transmit 300 bits and record them as ground truth. for comparison.

Metric. We set two metrics to quantify the effectiveness of DiskSpy: bit error rate (BER) [53] and signal-to-noise ratio (SNR) [15]. BER represents the accuracy of data transmission. It is the ratio of the number of incorrectly recognized bits to that of all received bits. SNR is defined by the ratio of the strength of the recovered vibration signal to that of background noise, which measures the quality of the recovered signal.

7.2 Overall Performance

In this section, we assess the bit transmission efficacy of the covert channel. To evaluate the performance of DiskSpy comprehensively, we also compare it with a benchmark (Tx bf), which only considers vibration SNR-guided Tx beamforming at the Tx end of radar without performing multi-Rx phase coherent integration at the Rx end. The BERs of nine HDDs are shown in Fig. 18. It can be seen that the average BER of DiskSpy is 0.65%, while that of the Tx bf method is as high as 8.54%. Hence, our multi-Rx phase coherent integration approach can effectively improve the data recovery performance. With these two strategies, Seagate BarraCuda 5400,2T has the smallest BER of 0.47%. Although Hitachi 7200,1T shows the worst performance, its BER is still less than 0.94%. This demonstrates that DiskSpy can accurately recover the information transmitted via HDD vibrations.

7.3 Distance and Orientation

Owing to the random deployment and mobility of air-gapped devices, the position of the HDD may not be fixed, e.g., varied distance and orientation. To evaluate the practicality of DiskSpy, we conduct two experiments with Seagate Pipeline 7200, 2T HDD under varying HDD distances and angles.


Distance. We first evaluate the impact of HDD distance on DiskSpy at a fixed bit rate, and then explore the tradeoff

Table 1: Tradeoff between HDD distance, bit rate, and BER.

Metrics	HDD Distance			
	2 m	10 m	18 m	26 m
Bit rate	20bps	20bps	16bps	2bps
BER	1.04%	1.15%	0.83%	0.87%
SNR	32.02dB	30.36dB	40.66dB	36.88dB

among HDD distance, bit rate, and BER. In the first experiment, we fix the HDD angle at 0° and vary the HDD distance from 2m to 22m with a step of 4m. In this experiment, the data transmission bit rate is set to the default 16bps. Fig. 19(a) shows the BERs and SNRs corresponding to each HDD distance. It can be observed that when the HDD distance increases from 2m to 22m, the BER only increases from 0.63% to 0.94%, while the SNR only drops from 51.01dB to 33.76dB. This indicates that the performance of DiskSpy will degrade as the HDD distance increases, while DiskSpy still performs well within 22m, i.e., BER<1% when the bit rate is 16bps. To further assess DiskSpy’s performance at a longer distance, we extend the HDD distance to 26m while maintaining the bit rate at 16bps. In this case, we find that the performance of DiskSpy will drop significantly, i.e., the BER>10%. This result implies that at a longer distance of 26m, DiskSpy cannot support high bit rate data transmission. Then, to demonstrate this issue, we explore the tradeoff among HDD distance, bit rate, and BER. Specifically, we measure the maximum bit rates at different HDD distances to ensure that the BER maintains within an acceptable range, i.e., BER<1.2%. Tab. 1 shows the results. We can observe that, as the HDD distance increases, the maximum bit rate that satisfies BER<1.2% gradually decreases. In particular, at the HDD distances of 2m and 10m, the maximum bit rate reaches 20bps. When the HDD distance is increased to 18m, the maximum bit rate drops to 16bps, and when the distance is further increased to 26m, the maximum bit rate plummets to 2bps. To investigate DiskSpy’s capability to handle higher bit rates at 26m, we further assess its performance at a bit rate of 5bps and find that the BER increases to 5.92%, exceeding the acceptable range. Furthermore, according to Fig. 19(a) and Tab. 1, we can also find that the SNR is affected by both HDD distance and bit rate. Specifically, as the bit rate increases from 16bps (Fig. 19(a)) to 20bps (Tab. 1), the SNRs for HDD distances of 2m and 10m decrease from 51dB and 48dB to 32.02dB and 30.36dB, respectively. These experimental results reveal that DiskSpy can achieve effective and accurate sensitive data transmission within 22m. Nevertheless, to launch attacks over a longer distance, attackers need to reduce the bit rate to retain the reliability of information transmission.

Orientation. In this experiment, we fix the HDD distance at 10m and vary the HDD angle from -60° to 60° in a step of 15° . Fig. 19(b) presents the BERs and SNRs corresponding to each HDD angle. It can be found that at all HDD angles, the BER remains below 0.8%, while the SNR consistently



(a) The setups of LOS and NLOS (b) DiskSpy’s performance in LOS and NLOS attack scenarios.

Figure 20: Performance in LOS and NLOS attack scenarios.

exceeds 42dB. Therefore, under a relatively long distance, the position of the HDD does not impact the attack effectiveness too much.

7.4 Robustness Study

We use Seagate Pipeline 7200, 2T HDD to assess DiskSpy’s NLOS performance and the impact of various environmental factors.

NLOS. In practice, obstacles like tables and cabinets may block the line-of-sight path between the radar and the air-gapped device. To investigate DiskSpy’s performance under NLOS conditions, we conduct comparative evaluations of its performance in both LOS and NLOS scenarios. Specifically, as shown in Fig. 20(a), we simulate the NLOS scenario and the LOS scenario by introducing or removing obstacles in the environment that block mmWave signals. In the LOS scenario, the beam is directly steered to the target along the LOS path. Conversely, in the NLOS scenario, the beam is steered to the target via reflections from the reflector in the environment. Then, as shown in Fig. 20(b), we vary the HDD distance R_0 to evaluate DiskSpy’s performance in different NLOS scenarios. It can be observed that the SNR consistently exceeds 46.8dB in both LOS and NLOS scenarios at various HDD distances. Additionally, the SNR values of DiskSpy in LOS and NLOS scenarios are very close. It indicates that DiskSpy performs well even in NLOS attack scenarios.

Environmental factors. To verify the robustness of our attack method in complex environments, we place various static and dynamic objects in the air-gapped environment and evaluate the DiskSpy performance. Specifically, we launch covert-channel attacks under a base environment (an office), with static objects surrounded, with vibrated smartphones interfered, with persons walking around, and with fans running, respectively. The resulting BERs and SNRs are depicted in Fig. 21. It can be found that the interfering components indeed lead to an increase in BER and a decrease in SNR. However, even a running fan close to the HDD causes only a slight performance degradation. In this case, BER is 0.71% and SNR is 44.43dB. Under the harshest conditions, i.e., with persons walking around, the BER is still as low as 0.78%, while the




Figure 21: Environmental factors.




Figure 22: HDD mounting manners.




Figure 23: Transmission Capacity.

SNR is higher than 41.07dB. This is because of the disparity in motion frequencies between environmental interference and the HDD. According to the HDD vibration frequency, DiskSpy can identify an optimal beam direction for vibration sensing utilizing the vibration SNR-guided Tx beamforming.

7.5 Impact of HDD Mounting Manner

There are four conventional mounting manners for HDDs (Fig. 17): external connection to the computer (M1), mounting on the chassis wall (M2), mounting in a removable HDD tray (M3), and attaching to the internal components of the chassis (M4). The external connection mounting manner M1 can be further categorized into two types: vertical placement (M1_v) and horizontal placement (M1_h). To understand the effect of different mounting manners, we mount a Hitachi 7200,2T HDD with these five manners to evaluate DiskSpy’s performance. The results in Fig. 22 demonstrate that the mounting manners have little impact on DiskSpy’s performance. The BERs in all mounting manners are lower than 0.71%, and the SNRs are higher than 44.43dB. This is because when the HDD is mounted inside the chassis, the vibrations generated by the HDD are transmitted through the internal structure of the chassis to the exterior. In this case, the mmWave radar senses the vibration of the chassis wall to conduct the covert-channel attack. However, the amplitude of the chassis vibrations is smaller than the HDD vibration, leading to a decrease in the SNR of the covert channel. However, with the help of the Tx-Rx co-optimization method, DiskSpy can effectively sense the subtle vibrations of the chassis and precisely decode them with a BER lower than 0.71%.

7.6 Transmission Capacity

Bit rate. By default, the time slot of T_1 (i.e., seek slot) is set to 100ms and T_0 is 25ms, resulting in a 16bps data rate. We set the duration of T_0 and T_1 unbalanced because of inertia. The HDD requires a rise time to reach its maximum vibration amplitude. This setting can ensure a higher data transmission speed. To study the transmission capacity of DiskSpy, we evaluate the BERs with varied seek slots (75, 100, 125, 150, and 175ms) under two distances (a short one of 2m and a long

one of 10m). A smaller time slot can provide a higher data rate, and the peak data rates for these five seek slots can be 20, 16, 13.3, 11.4, and 10bps, respectively. The experiment results in Fig. 23 indicate that the BERs decrease as seek slots increase. This is because a larger seek slot would make the idle and seek statuses more distinguishable. However, even when the bit rate is 20bps, the BER remains as low as 1.2%.

Bit length. We also investigate the impact of transmission bit length on DiskSpy. Specifically, we assess the performance of DiskSpy when transferring four different data lengths, including 100, 1000, 2000, and 4000 bits, under different HDD distances (2m and 10m) and with different bit rates (10bps and 16bps). The experiment results are displayed in Tab. 2. It can be observed that the BERs increase with the data length under all distance-bit rate combinations, while the magnitude of this increase is relatively small (<0.12%). Even when transmitting 4000-bit sensitive data, the BER is still smaller than 0.72%. These results indicate that DiskSpy is capable of sustaining long-data transmission.

7.7 Case Study

We conduct two case studies to evaluate the performance of DiskSpy in real attack scenarios: nature file transmission and outdoor-to-indoor attack.

Nature file transmission. In the first case study, we use DiskSpy to eavesdrop on three common files: text, image, and audio. The lengths of them are 120, 841, and 4546 bits, respectively. The experiment results in Fig. 24 demonstrate that the decoded text, image, and audio are almost identical to the original sensitive data. The error bits for text, image, and audio files are 0, 5, and 33, respectively. Therefore, DiskSpy’s performance, characterized by low BER (< 0.72%) during the transmission of natural files, demonstrates significant attack capabilities in real-world scenarios.

Outdoor-to-indoor attack. The outdoor-to-indoor attack setup is shown in Fig. 25(a) and (b). We randomly place a chassis equipped with a compromised HDD (i.e., the victim) in an air-gapped office. The HDD positions are indicated by dots in Fig. 25(a) and (b). In this experiment, DiskSpy has no prior knowledge of the victim’s location. Then, the mmWave

Table 2: Impact of bit length

Distance-bit rate combination	BER under different bit length			
	100 bits	1000 bits	2000 bits	4000 bits
2m, 10bps	0.57%	0.58%	0.63%	0.69%
2m, 16bps	0.64%	0.65%	0.67%	0.72%
10m, 10bps	0.62%	0.64%	0.65%	0.70%
10m, 16bps	0.65%	0.66%	0.70%	0.72%

radar (i.e., the attacker) is placed outside the office to attack the victim inside the office. As the LOS path is blocked by a concrete wall (with a thickness of 25 cm), the mmWave signal can reach the victim only by traversing through the office’s wooden door (this door is closed and locked) and being reflected by the surrounding wall. The signal propagation path is depicted in Fig. 25(b). The experimental results in Fig. 25(c) show that even under *outdoor-to-indoor*, *wall-reflecting*, and *HDD location-unknown* conditions, DiskSpy can still accurately distinguish 0/1 bit. The low BER (0.71%) and high SNR (45.1dB) indicate that DiskSpy offers reliable performance in real-world attack scenarios.

8 Mitigations

Large isolation zone. DiskSpy utilizes mmWave signals to establish the covert channel. The energy of mmWave signals rapidly attenuates when propagating through the air. Thus, as long as the isolation zone is sufficiently large, e.g., a circular area with a radius larger than 26m, it is difficult for DiskSpy to maintain a low BER. However, constructing such a large isolation zone is costly to realize in real-world scenarios.

Geofencing mmWave. Geofencing prevents mmWave signals from approaching the target HDD, thus disrupting the physical link of the covert channel. This can be achieved through two types of methods. (1) Implementing a metallic enclosure around the HDD without physical conduct with the computers to prevent vibration conduction. Metal shield is necessary since other materials could be penetrated by mmWave. (2) Painting isolation walls with electromagnetic shielding paints. These geofencing solutions incur high cost and are difficult to strictly enforce in practical settings.

Honeytrap HDDs. Users can deploy some deceptive HDDs in the air-gapped environment to make the malicious mmWave radar capture meaningless vibration signals. Yet, for the honeytrap defense to be effective, the number of employed honeytrap HDDs should be large enough, rendering substantial defense costs. Besides, as DiskSpy utilizes a customized header to mark data packets, it becomes relatively easy to distinguish the target HDD from honeytrap ones.

9 Related Work

Air-gapped Covert Channels. Air-gapped isolation is considered effective in protecting sensitive data in computers. However, the data on air-gapped computers could leak through





	Sensitive Data	Decoded Result	Data Size	Error Bits
Text File	bxi1svoyf@3xze?	bxi1svoyf@3xze?	120 bits	0 bits
Image File			841 bits	5 bits
Audio File	 “Hello”		4560 bits	33 bits

Figure 24: Nature file transmission.

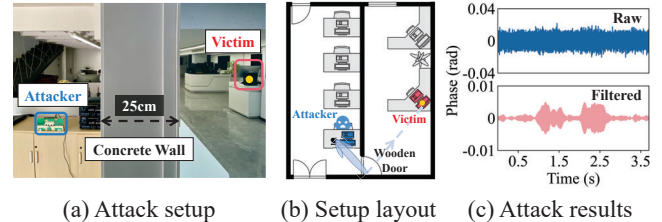


Figure 25: Outdoor-to-indoor attack.

covert-channel attacks, in which a variety of special media are exploited for information transmission, such as acoustic signal [18, 19, 37], RF leakage [15, 16], thermal emanation [4, 17], magnetic field [21], and contact vibration [13]. Table 3 compares DiskSpy with seven state-of-the-art air-gapped covert channels. These approaches are more or less limited in the following four aspects. (1) **LOS requirement.** In real attack scenarios, it is common for obstacles to block the LOS path, which poses challenges for those covert channels (such as directional antenna based [15] and thermal emanation-based [17]) that require the receiver to be precisely pointed to the infiltrated device. (2) **Efficiency.** A high bit rate can expedite the attack process. Yet, except for DiskSpy and the dedicated hardware-enabled method [15], the bit rates of the aforementioned attack strategies are relatively small (<15bps). (3) **Prior Knowledge.** Existing methods often require prior knowledge of target device information, such as location [13, 15, 17, 21, 31] and device parameters [43], when implementing covert-channel attacks. This requirement limits the practicality of these attack methods. (4) **Proximity.** Due to the signal attenuation, the effective attack range of most methods is limited to only a few meters [13, 17, 18, 21, 31]. Only two RF leakage based methods [15, 43] can support long-range (>20m) attacks. However, these two methods have made certain compromises. GSMem [15] requires expensive specialized hardware and EMLoRa [43] needs to sacrifice efficiency (i.e., bit rate) to achieve a long-range attack. DiskFiltration [18] also exploits the HDD to build a covert channel. It utilizes acoustic signals emitted from HDD to steal sensitive data. However, DiskFiltration is more susceptible to environmental noise interference, resulting in a low bit rate (3bps) over a short range (<2m). Compared with existing works, DiskSpy enables a practical covert-channel attack with **long**

Table 3: A Comparison with state-of-the-art Air-gapped Covert Channels.

System	Channel Modalities	Distance > 20m	Bit Rate > 15bps	NLOS	Prior Knowledge	Compromises for Realizing Long-range Attack
<i>Bitwhisper</i> [17]	Thermal	No	No	No	Location Information	No Long-range Support
<i>GSMem</i> [15]	RF Leakage	Yes	Yes	No	Location Information	Expensive Dedicated Hardware
<i>EMLoRa</i> [43]	RF Leakage	Yes	No	Yes	Target Device's Parameters	Sacrificing Bit Rate
<i>Odini</i> [21]	Magnetic	No	No	Yes	Location Information	No Long-range Support
<i>SpiralSpy</i> [31]	Tangential Velocity	No	No	Yes	Location Information	No Long-range Support
<i>Air-viber</i> [13]	Contact Vibration	No	No	No	Location Information	No Long-range Support
<i>DiskFiltration</i> [18]	Acoustic Signal	No	No	Yes	No	No Long-range Support
<i>DiskSpy</i>	Non-contact Vibration	Yes	Yes	Yes	No	No

attack distance and **high bit rate** even in **NLOS scenarios**.

mmWave-based Vibration Sensing. Compared with WiFi [3] and RFID [50], mmWave has the advantage of short wavelength, which gives it more opportunities to achieve high-precision vibration sensing. With the development of mmWave-based sensing, its precision has been improved from cm-level [35, 54] to μm -level [8, 9, 28, 52]. The fine-grained vibration sensing capability of mmWave enables a variety of novel applications. For example, mmVib [28] employs the mmWave signal to monitor the vibrations of industrial equipment for machinery health assessment. mmRipple [8] harnesses the mmWave radar to sense the modulated smartphone vibrations, thereby establishing a new communication channel. Additionally, various speech eavesdropping works [5, 10, 23, 24, 48] utilize the mmWave radar to sense the loudspeaker vibration to recover sound information. Yet, these mmWave-based vibration sensing methods primarily focus on high sensing precision at a limited sensing range (1-5m) in LOS scenarios [5, 8, 10, 23, 24, 48]. When LOS path is obstructed by impenetrable materials like metal (i.e., NLOS scenario), these methods become ineffective. Besides, the main reason that they can not support long-range vibration sensing is that, as the sensing distance increases, the SNR of the measured vibration signals decreases significantly, which directly affects the sensing precision, as illustrated in Sec. 5.2.1. Some works proposed SNR improvement methods to extend the sensing range but their performance is limited. For instance, mmEve [49] boosts the sensing distance by improving the SNR of mmWave signals via deep learning method. However, its maximum attack range is limited to 6-8 meters and it does not support NLOS scenarios. In the field of mmWave sensing, there are two prevalent techniques for extending sensing range: backscatter [45] and traditional beamforming [38]. However, they are primarily employed for long-range target localization and have their respective drawbacks when applied to our sensing tasks. The backscatter method requires the sensing target to be affixed with a dedicated tag. Apparently, this requirement is unrealistic in our attack scenario. Moreover, traditional beamforming is effective in improving the vibration signal SNR only when the sensing direction and the vibration direction coincide, as illus-

trated in Sec. 5.2.2. In this work, DiskSpy develops a novel SNR enhancement scheme, which for the first time, achieves μm -level vibration sensing at a long distance of 22m with COTS mmWave radars.

10 Discussion

Stealthiness analysis. To maintain stealthiness, DiskSpy utilizes seek operations to generate HDD vibrations for data encoding. Since seek operations are routine and frequently invoked HDD operations, it is difficult for victims to detect the covert-channel attack initiated by DiskSpy. Furthermore, to mitigate the impact on the victim's normal use of the HDD, we propose a data encoding scheme that divides sensitive data into multiple packets for transmission, thus allowing the transmission process to be interruptible. When the victim's processes require access to the HDD, DiskSpy prioritizes these tasks and records the index of the packet being transmitted. Once the victim's processes are complete, the transmission resumes from the recorded index. This strategy ensures seamless continuity of the information transmission process while minimizing the impact on the victim's normal HDD usage.

Initiator implantation methods. The size of sensitive data affects the data encoding time, thereby influencing the implantation methods. Specifically, if the size of sensitive data is small, such as a few kilobytes, the data encoding time is around several minutes. In this case, the malicious data encoding code can be inserted into some existing programs. Conversely, if the data size is substantial, such as several megabytes, the data encoding time may extend to dozens of hours. Under this circumstance, the malicious code needs to be implanted as dedicated malware into the air-gapped computer. These implantation methods have been extensively studied and can be realized via many attack methods [6, 30, 41].

11 Conclusion

This paper presents a new HDD-mmWave covert-channel attack, namely DiskSpy. It modulates the vibrations of HDDs to covertly send sensitive information from an air-gapped environment. Through remotely sensing the HDD vibration

statuses using mmWave signals, DiskSpy can obtain sensitive information on an air-gapped computer. To extend the attack distance, we devise a long-range subtle vibration sensing technique with the joint optimization at the Tx and Rx ends of mmWave radar. Extensive experiments show that DiskSpy is an efficient and robust long-range covert-channel attack.

Ethics Considerations

Throughout the experiment, we have carefully considered the ethical implications of this work. We conducted the experiments with the approval of the university's Institutional Review Board (IRB). We ensured that all experimentation was carried out in strict adherence to the necessary authorizations, avoiding any live testing on systems without appropriate approval. Additionally, any vulnerabilities identified during the research were promptly and responsibly disclosed to the relevant stakeholders, following best practices for ethical research conduct. These considerations align with ethical standards and foster responsible innovation in computer security and privacy.

Open Science

We are dedicated to openly sharing the research artifacts related to this work. We will make our system, DiskSpy¹, publicly available upon acceptance of the paper. The research artifacts will be released in a public repository, ensuring easy access and wide visibility under an open-source license. This commitment to open science is in line with the goal of promoting transparency and collaboration within the research community.

Acknowledgments

We sincerely thank our anonymous reviewers for their valuable feedback. This paper is supported by the National Natural Science Foundation of China under grants U21A20462 and 62372400, National Key R&D Program of China under grant No. 2023YFC3805602, "Pioneer" and "Leading Goose" R&D Program of Zhejiang under grant No. 2024C03287, Hong Kong GRF (Grant No. 15206123 and 15211924), and the Postdoctoral Fellowship Program of CPSF under grant No. GZC20241488.

References

- [1] An overview of social engineering malware: Trends, tactics, and implications. *Technology in Society*, 32(3):183–196, 2010.

¹The source code is available at <https://doi.org/10.5281/zenodo.14649224>

- [2] John C Bancroft. Introduction to matched filters. *CREWES Research*, 297, 2002.
- [3] Nan Bao, Jiajun Du, Chengyang Wu, Duo Hong, Junxin Chen, Robert Nowak, and Zhihan Lv. Wi-breath: A wifi-based contactless and real-time respiration monitoring scheme for remote healthcare. *IEEE Journal of Biomedical and Health Informatics*, 27(5):2276–2285, 2023.
- [4] Davide B. Bartolini, Philipp Miedl, and Lothar Thiele. On the capacity of thermal covert channels in multi-cores. In *Proceedings of the European Conference on Computer Systems (EuroSys)*, 2016.
- [5] Suryoday Basak and Mahanth Gowda. mmspy: Spying phone calls using mmwave radars. In *IEEE Symposium on Security and Privacy (SP)*, 2022.
- [6] Beyondtrust. What is a logic bomb? <https://www.beyondtrust.com/resources/glossary/logic-bomb>, 2023.
- [7] Ping Chen, Lieven Desmet, and Christophe Huygens. A study on advanced persistent threats. In *International Conference on Communications and Multimedia Security (CMS)*, 2014.
- [8] Kaiyan Cui, Qiang Yang, Yuanqing Zheng, and Jinsong Han. mmripple: Communicating with mmwave radars through smartphone vibration. In *Proceedings of the ACM International Conference on Information Processing in Sensor Networks (IPSN)*, 2023.
- [9] Lei Ding, Murtaza Ali, Sujeet Patole, and Anand Dabak. Vibration parameter estimation using fmcw radar. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016.
- [10] Yiwen Feng, Kai Zhang, Chuyu Wang, Lei Xie, Jingyi Ning, and Shijia Chen. mmeavesdropper: Signal augmentation-based directional eavesdropping with mmwave radar. In *IEEE Conference on Computer Communications (INFOCOM)*, 2023.
- [11] Benjamin Friedlander. On transmit beamforming for mimo radar. *IEEE Transactions on Aerospace and Electronic Systems*, 48(4):3376–3388, 2012.
- [12] Alexander S. Gillis. air gap (air gapping). <https://www.techtarget.com/whatis/definition/air-gapping>, 2022.
- [13] Mordechai Guri. Air-viber: Exfiltrating data from air-gapped computers via covert surface vibrations. *CoRR*, abs/2004.06195, 2020.

- [14] Mordechai Guri, Ofer Hasson, Gabi Kedma, and Yuval Elovici. Visisplot: An optical covert-channel to leak data through an air-gap. *CoRR*, abs/1607.03946, 2016.
- [15] Mordechai Guri, Assaf Kachlon, Ofer Hasson, Gabi Kedma, Yisroel Mirsky, and Yuval Elovici. Gsmem: Data exfiltration from air-gapped computers over GSM frequencies. In *Proceedings of the USENIX Security Symposium*, 2015.
- [16] Mordechai Guri, Gabi Kedma, Assaf Kachlon, and Yuval Elovici. Airhopper: Bridging the air-gap between isolated networks and mobile phones using radio frequencies. In *Proceedings of the International Conference on Malicious and Unwanted Software (MALWARE)*, 2014.
- [17] Mordechai Guri, Matan Monitz, Yisroel Mirski, and Yuval Elovici. Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations. In *Proceedings of the IEEE Computer Security Foundations Symposium (CSF)*, 2015.
- [18] Mordechai Guri, Yosef A. Solewicz, Andrey Daidakulov, and Yuval Elovici. Diskfiltration: Data exfiltration from speakerless air-gapped computers via covert hard drive noise. *CoRR*, abs/1608.03431, 2016.
- [19] Mordechai Guri, Yosef A. Solewicz, Andrey Daidakulov, and Yuval Elovici. Fansmitter: Acoustic data exfiltration from (speakerless) air-gapped computers. *CoRR*, abs/1606.05915, 2016.
- [20] Mordechai Guri, Boris Zadov, and Yuval Elovici. Led-it-go: Leaking (A lot of) data from air-gapped computers via the (small) hard drive LED. In *Proceedings of the Conference on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, 2017.
- [21] Mordechai Guri, Boris Zadov, and Yuval Elovici. Odini: Escaping sensitive data from faraday-caged, air-gapped computers via magnetic fields. *IEEE Transactions on Information Forensics and Security*, 15:1190–1203, 2020.
- [22] Thomas Horton King, Jizheng He, Chun-Kai Yao, Akarsh Prabhakara, Mohamad Alipour, Swarun Kumar, Anthony Rowe, and Elahe Soltanaghahi. Platypus: Sub-mm micro-displacement sensing with passive millimeter-wave tags as "phase carriers". In *Proceedings of the International Conference on Information Processing in Sensor Networks (IPSN)*, 2023.
- [23] Pengfei Hu, Wenhao Li, Riccardo Spolaor, and Xiuzhen Cheng. mmecho: A mmwave-based acoustic eavesdropping method. In *Proceedings of the ACM Turing Award Celebration Conference (TURC)*, 2023.
- [24] Pengfei Hu, Yifan Ma, Panneer Selvam Santhalingam, Parth H Pathak, and Xiuzhen Cheng. Milliear: Millimeter-wave acoustic eavesdropping with unconstrained vocabulary. In *IEEE Conference on Computer Communications (INFOCOM)*, 2022.
- [25] Imperva. What is a phishing attack. <https://www.imperva.com/learn/application-security/phishing-attack-scam/>, 2022.
- [26] Texas Instrument. Awr2243 single-chip 76- to 81-ghz fmcw transceiver. <https://www.ti.com/lit/ds/symlink/awr2243.pdf>, 2022.
- [27] Texas Instrument. Mmwcas-dsp-evm. <https://www.ti.com/tool/MMWCAS-DSP-EVM>, 2023.
- [28] Chengkun Jiang, Junchen Guo, Yuan He, Meng Jin, Shuai Li, and Yunhao Liu. mmvib: micrometer-level vibration measurement with mmwave radar. In *Proceedings of the ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2020.
- [29] Brett Kelly. Everything you need to know about hard drive vibration. <https://www.ept.ca/features/everything-need-know-hard-drive-vibration/>, 2016.
- [30] Paul K Kerr, John Rollins, and Catherine A Theohary. *The stuxnet computer worm: Harbinger of an emerging warfare capability*. Congressional Research Service, 2010.
- [31] Zhengxiong Li, Baicheng Chen, Xingyu Chen, Huining Li, Chenhan Xu, Feng Lin, Chris Xiaoxuan Lu, Kui Ren, and Wenyao Xu. Spiralspy: Exploring a stealthy and practical covert channel to attack air-gapped computing devices via mmwave sensing. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2022.
- [32] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, et al. Meltdown: Reading kernel memory from user space. *Communications of the ACM*, 63(6):46–56, 2020.
- [33] Anna Maricheva. What is air gap and how does it impact your cybersecurity? <https://softteco.com/blog/what-is-air-gap>, 202.
- [34] Adriano Meta, Peter Hoogeboom, and Leo P Ligthart. Signal processing for fmcw sar. *IEEE Transactions on Geoscience and Remote Sensing*, 45(11):3519–3532, 2007.
- [35] Ilya Mikhelson, Sasan Bakhtiari, Thomas W. Elmer, and Alan V. Sahakian. Remote sensing of heart rate and

patterns of respiration on a stationary subject using 94-ghz millimeter-wave interferometry. *IEEE Transactions on Biomedical Engineering*, 58(6):1671–1677, 2011.

- [36] Savita Mohurle and Manisha Patil. A brief study of wannacry threat: Ransomware attack 2017. *International journal of advanced research in computer science*, 8(5):1938–1940, 2017.
- [37] Rajalakshmi Nandakumar, Alex Takakuwa, Tadayoshi Kohno, and Shyamnath Gollakota. Covertband: Activity information leakage using music. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(3):87:1–87:24, 2017.
- [38] Clemens Pfeffer, Reinhard Feger, Christoph Wagner, and Andreas Stelzer. Fmcw mimo radar system for frequency-division multiple tx-beamforming. *IEEE Transactions on Microwave Theory and Techniques*, 61(12):4262–4274, 2013.
- [39] Sandeep Rao. Introduction to mmwave sensing: Fmcw radars. *Texas Instruments (TI) mmWave Training Series*, 2017.
- [40] Mark A Richards et al. *Fundamentals of radar signal processing*. 2005.
- [41] Fatima Salahdine and Naima Kaabouch. Social engineering attacks: A survey. *Future internet*, 11(4):89, 2019.
- [42] Stefan Schneegass, Alia Saad, Roman Heger, Sarah Delgado Rodriguez, Romina Poguntke, and Florian Alt. An investigation of shoulder surfing attacks on touch-based unlock events. In *Proceedings of the ACM on Human-Computer Interaction (HCI)*, 2022.
- [43] Cheng Shen, Tian Liu, Jun Huang, and Rui Tan. When lora meets emr: Electromagnetic covert channels can be super resilient. In *IEEE Symposium on Security and Privacy (SP)*, 2021.
- [44] Cong Shi, Tianfang Zhang, Zhaoyi Xu, Shuping Li, Donglin Gao, Changming Li, Athina Petropulu, Chung-Tse Michael Wu, and Yingying Chen. Privacy leakage via speech-induced vibrations on room objects through remote sensing based on phased-mimo. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2023.
- [45] Elahe Soltanaghaei, Akarsh Prabhakara, Artur Balanuta, Matthew Anderson, Jan M Rabaey, Swarun Kumar, and Anthony Rowe. Millimetro: mmwave retro-reflective tags for accurate, long range localization. In *Proceedings of the ACM Conference on Mobile Computing and Networking (MobiCom)*, 2021.
- [46] Gaurav Somani, Manoj Singh Gaur, Dheeraj Sanghi, Mauro Conti, and Rajkumar Buyya. Ddos attacks in cloud computing: Issues, taxonomy, and future directions. *Computer Communications*, 107:30–48, 2017.
- [47] Naresh Tandon, V.V.P. Rao, and V. P. Agrawal. Vibration and noise analysis of computer hard disk drives. *Measurement*, 39:16–25, 2006.
- [48] Chao Wang, Feng Lin, Tiantian Liu, Ziwei Liu, Yijie Shen, Zhongjie Ba, Li Lu, Wenyao Xu, and Kui Ren. mmphone: Acoustic eavesdropping on loudspeakers via mmwave-characterized piezoelectric effect. In *IEEE Conference on Computer Communications (INFOCOM)*, 2022.
- [49] Chao Wang, Feng Lin, Tiantian Liu, Kaidi Zheng, Zhibo Wang, Zhengxiong Li, Ming-Chun Huang, Wenyao Xu, and Kui Ren. mmeve: eavesdropping on smartphone’s earpiece via cots mmwave device. In *Proceedings of the ACM Conference on Mobile Computing and Networking (MobiCom)*, 2022.
- [50] Lei Yang, Yao Li, Qiongzhen Lin, Huanyu Jia, Xiang-Yang Li, and Yunhao Liu. Tagbeat: Sensing mechanical vibration period with COTS RFID systems. *IEEE/ACM Transactions on Networking*, 25(6):3823–3835, 2017.
- [51] Qian Yang, Hengxin Wu, Qianyi Huang, Jin Zhang, Hao Chen, Weichao Li, Xiaofeng Tao, and Qian Zhang. Side-lobe can know more: Towards simultaneous communication and sensing for mmwave. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(4):1–34, 2023.
- [52] Yanni Yang, Huafeng Xu, Qianyi Chen, Jiannong Cao, and Yanwen Wang. Multi-vib: Precise multi-point vibration monitoring using mmwave radar. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(192):1–26, 2023.
- [53] Zhice Yang, Qianyi Huang, and Qian Zhang. Nicscatter: Backscatter as a covert channel in mobile devices. In *Proceedings of the ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2017.
- [54] Zhicheng Yang, Parth H. Pathak, Yunze Zeng, Xixi Liran, and Prasant Mohapatra. Vital sign and sleep monitoring using millimeter wave. *ACM Transactions on Sensor Networks*, 13(2):14:1–14:32, 2017.

A Data Encoding algorithm

Algorithm 1: Data Encoding

Input: D : Sensitive data, $flag_{end}$: Encoding end indicator, $index_p$: Packet index being transmitted, $length_p$: Packet length, $header_p$: Packet header, T_0 : The duration to transmit '0', T_1 : The duration to transmit '1', $address_s$: The source address, $address_d$: The destination address


```

1 Function dataEncoding( $D, index_p$ ):
2    $index_d = index_p \times length_p$ 
3   while  $index_d < len(D)$  do
4     if  $index_d + length_p < len(D)$  then
5        $payload = D[index_d : index_d + length_p]$ 
6     else
7        $payload = zeroPad(D[index_d : len(D)])$ 
8     end
9      $packet = [header_p, payload]$ 
10    bitEncoding( $packet$ )
11     $index_d = index_d + length_p$ 
12     $index_p = index_p + 1$ 
13  end
14  bitEncoding( $flag_{end}$ )
15 Function bitEncoding( $Packet$ ):
16   $seek(address_s)$ 
17   $address_{cur} = address_s, address_{next} = address_d$ 
18  for  $b$  in  $Packet$  do
19    if  $b = '0'$  then
20       $sleep(T_0)$ 
21    end
22    if  $b = '1'$  then
23      for  $T_1$  do
24         $seek(address_{next})$ 
25         $swap(address_{cur}, address_{next})$ 
26      end
27    end
28  end

```

B Evaluation on optimal beam identification

DiskSpy identifies the optimal beam direction through vibration SNR-guided Tx beamforming. To evaluate its effectiveness, we first define a metric angle error (AE) which is quantified by the absolute difference between the theoretically optimal beam direction and the best beam direction identified by vibration SNR-guided Tx beamforming. Then, we conduct an experiment under different HDD deflection angles to measure the AEs of the optimal beam direction identification, as shown in Fig. 26(a). Specifically, we alter the deflection angle of HDD from -30° to 30° at intervals of 10° . For each deflection angle, we will calculate the theoretical optimal beam direction at that time as a ground truth. Then, we compare the best beam direction identified by our approach with the ground truth to measure the performance of the proposed vibration SNR-guided Tx beamforming. Fig. 26(b) presents the AEs corresponding to each HDD deflection angle. It can be observed that the AEs of all deflection situations are lower than 2° . Thus, our proposed Tx beamforming approach per-



(a) Altering HDD deflection and calculating the theoretical optimal beam direction. (b) AEs under different HDD deflection angles.

Figure 26: Evaluations on the performance of optimal beam direction identification.

forms well in steering Tx beam to the target HDD.

C μm -level vibration sensing with mmWave

According to Sec. 3.2, in an ideal scenario (e.g., no signal attenuation and no hardware imperfection), we can accurately measure the vibrating target's micro-displacement (ΔR) by tracking the phase changes ($\Delta\phi_{ideal}$) of the corresponding range bin:

$$\Delta\phi_{ideal} = \frac{4\pi f_c \Delta R}{c}. \quad (5)$$

However, in the realistic scenario, the received mmWave signals are significantly affected by noises [22] such as free-space path loss, and phase offsets during the mixing process and down-chirping. Therefore, the precision of target's displacement derived from phase changes ($\Delta\phi_{realistic}$) is affected by these noises:

$$\Delta\phi_{realistic} = \frac{4\pi f_c \Delta R}{c} + \phi_{noise}, \quad (6)$$

where ϕ_{noise} is the noise in phase values. According to [44], the value of ϕ_{noise} impacts the resolution of measured displacement. In other words, a displacement can be detected and observed only when the phase change induced by the displacement is larger than ϕ_{noise} :

$$\frac{4\pi f_c \Delta R}{c} > \phi_{noise}. \quad (7)$$

Therefore, the target's displacement resolution can be calculated as $\Delta R_{resolution} = \frac{c\phi_{noise}}{4\pi f_c}$. Taking the case where the attack distance is 22m (Fig. 14(c)) as an example, the value of ϕ_{noise} is measured as 4.933×10^{-3} rads. In this case, given a $f_c = 77\text{GHz}$ mmWave radar, the target's displacement resolution $\Delta R_{resolution} = 1.49\mu\text{m}$. Hence, the measured displacement resolution is sufficiently high to detect μm -level HDD vibrations at a long range.