ODINI: Escaping Sensitive Data From Faraday-Caged, Air-Gapped Computers via Magnetic Fields

Mordechai Guri, Boris Zadov, and Yuval Elovici

***Abstract*— Air-gapped computers are devices that are kept isolated from the Internet, because they store and process sensitive information. When highly sensitive data is involved, an air-gapped computer might also be kept secluded in a Faraday cage. The Faraday cage prevents the leakage of electromagnetic signals emanating from various computer parts, which may be picked up remotely by an eavesdropping adversary. The air-gap separation, coupled with the Faraday shield, provides a high level of isolation, preventing the potential leakage of sensitive data from the system. In this paper, we show how attackers can bypass Faraday cages and air-gaps in order to leak data from highly secure computers. Our method is based on exploitation of the magnetic field generated by the computer's CPU. Unlike electromagnetic radiation (EMR), low frequency magnetic fields propagate through the air, penetrating metal shielding such as Faraday cages (e.g., a compass still works inside a Faraday cage). Since the CPU is an essential part of any computer, the magnetic covert channel is relevant to virtually any device with a CPU: desktop PCs, servers, laptops, embedded systems, and Internet of Things (IoT) devices. We introduce a malware codenamed 'ODINI' that can control the low frequency magnetic fields emitted from the infected computer by regulating the load of the CPU cores. Arbitrary data can be modulated and transmitted on top of the magnetic emission and received by a magnetic 'bug' located nearby. We implement a malware prototype and discuss the design considerations along with the implementation details. We also show that the malicious code does not require special privileges (e.g., root) and can successfully operate from within isolated virtual machines (VMs) as well. Finally, we propose different types of defensive countermeasures such as signal detection and signal jamming to cope with this type of threat (demonstration video: https://www.youtube.com/watch?v=h07iXD-aSCA).**

***Index Terms*— Network security, air gaps, computer viruses.**

I. Introduction

ONE of the main goals of advanced persistent threat (APT) attacks is to steal sensitive information from compromised organizations. Currently, defending computer networks from APTs and sophisticated cyber-attacks is a complicated task, which involves maintaining multiple layers of security systems. This includes updating protection software on the host computers, configuring firewalls and routers, managing access controls, using centralized credential systems, and so on. Nevertheless, despite a high degree of protection, as long as the local area network is connected to the Internet, a motivated adversary will find a way to breach the network, evade security mechanisms, access sensitive data, and transfer it outside to the attacker [1]–[3].

In order to provide protection from such breaches, organizations often store their sensitive data on air-gapped networks. In this approach, any type of physical or logical connection between the local network and the Internet is strictly banned. Air-gapped networks are widely used in military and defense systems, critical infrastructure, the finance sector, and other industries [4], [5]. Two famous examples of air-gapped networks are the NSANet and the Joint Worldwide Intelligence Communications System (JWICS), classified (top secret) networks belonging to the United States' Defense Intelligence Agency [6].

A. Air-Gap Infiltration

Despite their isolation, air-gapped networks are not immune to breaches. It has been shown that attackers can successfully penetrate air-gapped network by using complex attack vectors, such as supply chain attacks, malicious insiders, and social engineering [7], [8]. Note that many modern APTs are capable of infecting air-gapped systems, including ProjectSauron [9], Fanny [10], Regin [11], Stuxnet [12], and Agent.BTZ [13]. HammerDrill and Brutal Kangaroo [14], [15], disclosed by WikiLeaks in 2017, are two types of attacking frameworks that can use removable media as a covert channel to compromise air-gapped systems [14]. Using these tools, attackers can bypass security systems and gain a foothold within air-gapped networks.

B. Air-Gap Covert Channels

After installing a malware in the air-gapped network, the attacker may, at some point, wish to retrieve sensitive information such as files, encryption keys, passwords, and more. However, despite the fact that *infiltration* of air-gapped systems has been shown feasible, the *exfiltration* of data from air-gapped system remains a challenge. Over the years, various out-of-band communication methods to leak data through air-gaps have been proposed. In particular, electromagnetic based

Manuscript received May 24, 2018; revised July 9, 2019; accepted August 15, 2019. Date of publication August 29, 2019; date of current version December 11, 2019. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Jean-Luc Danger. *(Corresponding author: Mordechai Guri.)*

The authors are with the Department of Information Systems Engineering, Ben Gurion University of the Negev, Be'er Sheva 8410501, Israel (e-mail: gurim@post.bgu.ac.il).

Digital Object Identifier 10.1109/TIFS.2019.2938404

covert channels have been studied by the academic research community for at least two decades. In this type of communication, a malware exploits the electromagnetic waves radiating from computer components to modulate binary information. The radiating components might be communication cables, computer buses, and hardware peripherals [16]–[20]. Such a technique is referred to in NSA jargon as a 'TEAPOT' attack [21].

C. Faraday Shielding

To cope with electromagnetic leakage, highly sensitive equipment might be placed within metal enclosures known as Faraday shielding or Faraday cages. A Faraday cage is made of conducting material (e.g., wire mesh or metal plates) that shields the area inside the cage from external electric fields. In the context of protecting sensitive equipment, Faraday shields are used to block electromagnetic waves from 1) being leaked from the shielded area or 2) penetrating into it. The most simple case of Faraday shielding is when it is implemented in the computer cabling (e.g., Ethernet, USB, and HDMI cables) to limit their electromagnetic emissions [22]. A Faraday shield is a type of small enclosure that can be deployed to protect entire systems such as desktop PCs and display screens [23], [24], but they may also be used to protect entire rooms and even buildings [25]. Faraday shielding renders most air-gap covert channels ineffective, since it prevents the leakage of electromagnetic signals outside to the attacker.

D. Our Contribution

In this paper, we present a new type of covert channel that can be used to exfiltrate information from air-gapped computers through Faraday cages. Our method uses low frequency magnetic fields generated by a computer's CPU. These fields penetrate metal shields, and hence can be used to bypass the protective Faraday cages.

The following aspects/points represent the contributions of our paper

- **Leaking through Faraday shielding.** We introduce a covert channel that can evade Faraday isolation. That is, it can work in highly secured systems which are kept within Faraday cages where other types of electromagnetic covert channels fail. As far as we know, this is the first work that discusses the topic of Faraday cages and their evasion using covert channels.
- **Air-gap covert channel.** The communication channel we introduce is an air-gap covert channel. That is, regardless of its ability to bypass Faraday shielding, it is capable of leaking data from disconnected, air-gapped computers.
- **Bypassing virtual machine (VM) isolation.** Virtual machines are often used as a security measure to add a layer of isolation between the VM and the external environment. We show that the covert channel works even when the malicious code is executed on 'virtualized' hardware, in an isolated VM.
- **Evaluation of a magnetic 'bug.'** We introduce the concept of a maliciously implanted magnetic receiver ('bug'), similar to microphone bugs and radio frequency (RF) receivers used with traditional covert channels.

The paper is structured as follow: In Section II we present related work. Section III describes the attack model. Scientific background on magnetic fields and Faraday cages is provided in Section IV. The modulation, signal generation, data encoding, and transmission protocols are described in Section V. In Section VI we present the analysis and evaluation. Countermeasures are discussed in Section VII, and we present our conclusions in Section VIII.

II. Related Work

Conventional covert channels assume the existence of network connectivity between the attacker and the target network. These types of covert channels have been widely studied and discussed in prior academic works and literature [26], [27]. Using covert channels, attackers may hide data within legitimate network traffic (e.g., HTTPS, FTP, and DNS), conceal it in images (steganography), or encode it in packet timings [26], [28]–[30]. In cases where there is no direct connection with the target network, the attacker may resort to so-called 'air-gap' or 'out-of-band' covert channels. Guri and Elovici present a taxonomy of the class of malware that exploits air-gap covert channels in order to bridge the air-gap between isolated networks and attackers [21]. Carrara and Adams provide a thoughtful survey of these covert channels [31]. They are classified into electromagnetic, acoustic, optical, thermal, magnetic and seismic covert channels.

A. Electromagnetic

The electromagnetic based covert channel has been the most researched topic in this field for at least twenty years. Kuhn and Anderson showed how attackers can leak data from air-gapped computers by controlling the electromagnetic waves emanating from display screens [17]. In their method, called 'soft tempest,' a malicious code encodes information over AM signals generated by certain bitmap patterns displayed on the screen. Based on this work, in 2001 Thiele [32] presented an open-source program dubbed 'Tempest for Eliza,' which uses the computer monitor to transmit AM radio signals; the transmissions can be heard from a nearby simple radio receiver. In 2014, Guri et al introduced AirHopper, malware that can exfiltrate data from air-gapped networks to nearby mobile phones using controllable electromagnetic signals in the FM radio band emanating from the video cable [4], [16]. Later on, in 2015, Guri et al presented GSMem, malware that leaks data from air-gapped computers using frequencies in the cellular band emitted from memory buses [20]. In their method, they use a multichannel memory architecture to amplify the transmission power. The transmission is then received by a rootkit placed on baseband firmware of a compromised mobile phone. Researchers also proposed using USB data bus and GPIO ports to generate covert electromagnetic signals for data exfiltration [33], [34].

B. Acoustic

In acoustic covert channels, data is transmitted via audible or ultrasonic sound waves. Audio communication between

computers was first reviewed by Madhavapeddy et al in 2005 [35]. Later on, research discussed using ultrasonic sound waves (18-22kHz) to transmit data between air-gapped laptops using their speakers and microphones [36], [37]. However, the aforementioned acoustic and ultrasonic methods are not relevant when speakers or microphones are not installed in the computer. In 2016, Guri et al presented Fansmitter [38] and DiskFiltration [39], two methods enabling exfiltration of data via sound waves when the computers are not equipped with speakers or audio hardware; the binary data is modulated via noise emitted from computer fans and the hard disk drive actuator arm.

C. Optical

Several studies have proposed the use of optical emanation for covert communication. Loughry and Umphress proposed a malicious code that exfiltrates data by blinking the Caps Lock, Num Lock, and Scroll Lock LEDs on the PC keyboard [40]. More recently, researchers presented covert channels that uses the hard drive indicator LED [41] and the router LEDs [42] in order to leak data from air-gapped networks. VisiSploit [43] is another optical based covert channel in which data is leaked through fast blinking images or low contrast bitmaps projected on the computer screen. Lopes and Aranha [44] presented a covert channel based on signals transmitted from IR LEDs in external USB devices attached to the air-gapped computer. In 2017, researchers presented a method that uses the IR LEDs present in surveillance and security cameras to exfiltrate and infiltrate air-gapped networks remotely [45].

D. Thermal

In 2015, Guri et al presented a thermal based method called BitWhisper [46]. In this technique, an attacker can establish bidirectional communication between two adjacent air-gapped computers using heat emissions. The heat is generated by CPU/GPU cores and received by thermal sensors that exist in the PC motherboard.

Magnetic communication in general is a known topic of research [47]. For example, the MagneLink Magnetic Communication System (MCS) is a system which provides through-the-earth emergency wireless communication based on magnetic fields [48]. Near-field magnetic induction (NFMI) communication is another type of magnetic method that allows short range communication between devices [49]. However, these types of communication methods require dedicated magnetic transmitters and receivers, which are not available in the case of our covert channel.

In the context of covert channels, Matyunin uses hard disk drives' (HDD) read/write operations to generate magnetic emissions, which can be measured by a nearby smartphone's magnetic sensor [50]. The smartphone needs to be located a few centimeters from a transmitting laptop, and the bitrate varies from 0.067 bit/sec to 2 bit/sec. However, their attack does not work on standard desktop workstations, since the generated signal is too weak (the smartphone must be placed *on* the workstation chassis directly above the HDD location, in order to measure the signal). Our method differs from previous work ([50]) in the following respects:

1) **Air-gaps and Faraday shielding.** Our discussion and evaluation focus on a covert channel that is relevant to air-gap and Faraday isolation. To the best of our knowledge, this is the first paper that discusses the topic of Faraday cages evasion in the context of covert channels and cyber security measures.
2) **Signal strength.** The magnetic fields generated by our method are ten times stronger than the magnetic fields generated by read/write operations on HDDs. Consequently, our covert channel enables higher bitrates and transmission from greater distances.
3) **Hardware availability.** We propose generating magnetic fields via the CPU which is available on virtually any computerized device today, including devices without magnetic HDDs (e.g., SSDs, embedded and IoT devices, and so on).
4) **Signal control.** Our method enables us to control the frequency of the transmissions independently for each CPU core, and hence to be able to use more complex, faster modulation schemes.
5) **Stealth and evasion.** Our transmitting code can be executed from any user mode process. It uses basic CPU instructions, and does not perform HDD I/O operations (read/write). This makes it difficult for anomaly detection systems to identify the malicious activity of the transmitter.

III. Attack Model

The adversarial attack model requires running a malicious code on the targeted computer. In addition, there must also be a magnetic receiver hidden or positioned near the targeted system (this could take the form of the attacker or an insider physically carrying the receiver near the targeted system). The attack itself consists of two phases. A preliminary phase which includes the system infection and an active phase which includes the data exfiltration.

A. System Infection

In the initial phase, the attacker infects the target system or network with malware. As discussed, infecting highly secure and even air-gapped networks has been proven feasible. Note that several APTs discovered in recent years are capable of infecting air-gapped networks [10], [51]–[53]. As part of the targeted attack, the adversary may infiltrate the air-gapped networks using social engineering, supply chain attacks, or malicious insiders. The magnetic receiver can be implanted in close proximity to the targeted system, outside the Faraday shield. Another option is to physically carry the magnetic receiver near the targeted system temporarily, in a so called 'evil maid' attack [54].

B. Data Exfiltration

Having a foothold in the system, the malware starts retrieving interesting data for the attacker. The data might be textual

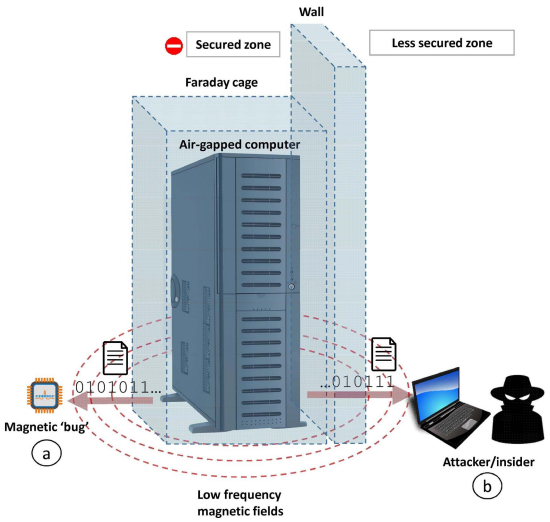

Fig. 1. An illustration of the magnetic covert channel (ODINI). Sensitive data is exfiltrated from the secured system, through air-gap and Faraday shielding. It is received by (a) a magnetic receiver planted nearby or (b) a magnetic receiver carried by an attacker or insider.

information, keystrokes logging, encryption keys, credential tokens, and more. the malware encode the data and transmit it via low frequency magnetic fields generated by the CPU. The leaked data is received by the nearby magnetic receiver. In the case of an implanted receiver, the data is sent to the attacker encrypted via the Internet (e.g., with Wi-Fi).

Note that although the described attack model is complicated, it is not beyond the capability of motivated and capable attackers. As a reward for their efforts, attackers can get their hands on valuable information, which is highly protected against other types of covert channels.

Figure 1 illustrates the magnetic covert channel described above. The data is received by (a) a magnetic receiver planted near the computer in the same room or (b) a magnetic receiver carried by an attacker/insider. In this case the attacker/insider is carrying the magnetic receiver in a less restricted zone, behind a wall.

IV. Scientific Background

In this section, we provide the scientific background necessary to understand the magnetic covert channel. We briefly introduce the concept of magnetic fields and discuss Faraday shielding.

A. Magnetic Field

Magnetic fields are produced when current flows in a straight wire and are propagated in space at the speed of light. A magnetic field at a given point is specified by its direction and strength and is mathematically represented by a vector field. The international system unit of the intensity for magnetic fields is the tesla (T). One tesla is defined as the field intensity generating one newton (N) of force per ampere (A) of current per meter of conductor. In practice, a magnetic field of one tesla is very strong and magnetic fields are commonly measured in units of milliteslas ($1mT = 10^{-3}T$) or microteslas ($1\mu T = 10^{-6}T$). Ampere's Law shows that the strength of the magnetic field around an electric current is proportional to the electric current. The strength of the magnetic field is proportional to the third power of the distance from the center of the wire [55]. More specifically, the magnetic flux density equation shows that the magnetic field's rapid decay is proportional to the inverse of the third power of the distance from the source:

$$B(r) = \nabla \times A = \frac{\mu_0}{4\pi}\left(\frac{3r(m \cdot r)}{|r|^5} - \frac{m}{|r|^3}\right) \quad (1)$$

where B is the strength of the magnetic field in teslas, and r is a distance from the source. The other parameters are the magnetic potential (A), magnetic permeability (μ_0), and the magnetic moment (m). Note that a scientific overview of the magnetic flux density equation is out of the scope of this paper, and we refer the interested reader to textbooks focusing on electromagnetics [56]. As can be seen in Eq. 1, the main disadvantage of the magnetic field is its rapid decay, which limits the distance of magnetic communication compared to that of electromagnetic communication [56]. In practice, magnetic fields are mostly used for the establishment of short range wireless communication between nearby devices, a technique commonly referred to as near-field magnetic induction communication [49].

B. Faraday Shielding

Faraday shielding is an enclosure used to block electromagnetic fields (e.g., radio transmissions) from leaking out or entering into the shielded system. From a scientific point of view, a Faraday shield is a case, which conducts all electromagnetic radiation on its surface. It makes the entire surface to be with an equal potential and prevents potential changes inside the enclosure. Faraday shields may be small in size when protecting computer systems, or very large when protecting entire rooms and laboratories [25]. Faraday shielding plays an important role in the field of emission security (EMSEC), particularly by providing protection from TEMPEST attacks. In this type of attack, adversaries intercept the electromagnetic radiation emanating from electronic equipment and reconstruct the information processed in the device [18]. Faraday shielding copes with this threat by preventing the leakage of electromagnetic signals from the shielded area. Generally, the shielding involves encompassing the device in a Faraday cage that does not permit stray electromagnetic emanations. It should also be noted that there are governmental and commercial standards (e.g., NATO SDIP-27 and NSTISSAM) which require limiting such emanation from devices for security and safety purposes [57].

C. Magnetic Fields and Metal Shielding

The propagation of electromagnetic and magnetic radiation in conducting mediums such as concrete is better at low frequencies [55], [58]. However, in the case of electromagnetic waves, the antenna required for low frequency transmissions is extremely long, since it is proportional to the wavelength. For example, an efficient transmission of an electromagnetic

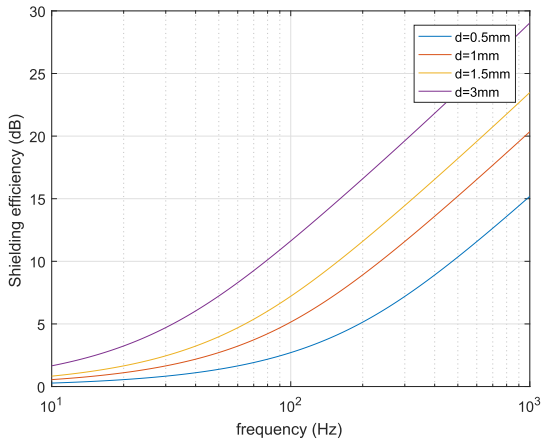

Fig. 2. Shielding efficiency of the closed metal shell at a thickness of 0.5 mm to 3 mm.

signal at 100kHz would require an antenna that is more than three kilometers long. Magnetic waves, on the other hand, do not depend on antenna length, and hence provide a practical alternative for wireless communication at low frequencies. Specifically, low frequency magnetic waves can propagate through dense medium such as metal, concrete, and soil [48], [55]. In the proposed covert channel, we generate magnetic fields at frequencies lower than 50Hz. It is known that very low frequency magnetic fields have low impedance and are difficult to block with metal shielding, since this would require very thick metal surfaces [59]. In Figure 2 we show the attenuation of a magnetic field (the reduction in magnitude of magnetic field strength) given a cubic metal shield, based on the shielding approximation formulas [60]. The field attenuation is measured in decibels (dB) and is equal to $20\ log\frac{E1}{E2}$, where E1 is the field intensity generated on one side of the shield, and E2 is the field intensity received on the other side of the shield. We calculate the efficiency of cubic metal cases at thicknesses of 0.5mm to 3mm in blocking magnetic fields at frequencies below 1000Hz. As can be seen, even thick metal shields are not efficient for very low frequencies ($<$ 50Hz), as the magnetic attenuation is 5dB at most.

V. Modulation

In this section we describe the signal generation algorithm and present the data modulation and the transmission protocol.

A. Signal Generation

As described in the scientific background section, moving charges in a wire generates a magnetic field. The magnetic field changes according to the acceleration of the charges in the wire. In a standard computer, the magnetic emanation stems primarily from wires that supply electricity from the main power supply to the motherboard. Since modern CPUs are energy efficient, the momentary workload of the CPU directly affects the dynamic changes in its power consumption [61]. By regulating the workload of the CPU, it is possible to govern its power consumption, and hence control the magnetic field generated. In the most elementary case, overloading the CPU with calculations will consume more current and consequentially will generate a stronger magnetic field. Intentionally starting and stopping the CPU workload allows us to generate a magnetic field at the required frequency and modulate binary data over it.

In our approach, the workload of each of the CPU core is managed independently from the other cores. Regulating the workload of each core separately enables greater control of the magnetic field generated. This approach has the following advantages:

1) **Using available cores.** Working at the resolution of cores allows us to use only the available, non-utilized cores. This way the covert channel won't interrupt other processes in the system.
2) **Controlling the signal strength.** By using different numbers of cores for the transmission, we can control the strength of the magnetic field. This allows us to employ amplitude modulations in which data is encoded on the amplitude level of the signal.
3) **Using multiple frequencies.** By controlling the workload of each core separately, we can use a different sub-carrier for each transmitting core. This allows us to employ a more efficient modulation scheme such as orthogonal frequency-division multiplexing (OFDM).

We generate a carrier wave at frequency f in one or more cores, by controlling the utilization of the CPU at a frequency correlated to f. To that end, we create worker threads and bound each thread to a specific core. The basic operation of a worker thread is described in algorithms 1-4.

Algorithm 1 Transmit $(nWorkers, vector, f, nCycles)$

```
for i ← 0 to nWorkers do
    mutex[i] = createMutex()
    acquiredMutex(mutex[i])
    createThread(workerThread, mutex[i], vector, f,
        nCycles)
end for
for j ← 0 to nWorkers do
    releaseMutex(mutex[i])
end for
```

Algorithm 2 Workerthread $(threadMutex, f, nCycles)$

```
bindThreadToCore()
acquire(threadMutex)
for i ← 0 to vector.length do
    signal(vector[i], f, nCycles)
    i++
end for
release(threadMutex)
```

The *transmit* function receives the number of worker threads to initiate $(nWorkers)$, the stream of bits to transmit $(vector)$, the frequency of the carrier signal (f), and the number of carrier wave cycles per bit $(nCycles)$. The transmitter creates a mutex object for each worker thread. These mutex objects are used to synchronize the worker threads at the beginning of the

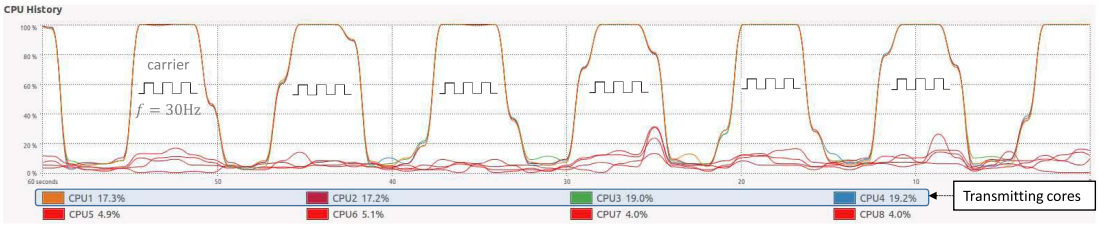

Fig. 3. The transmitter run in a eight-cores PC (CPU utilization view) with $f = 30\ Hz$. The transmitting threads are bound to cores CPU1-CPU4. Note that the CPU history graph presents a *moving average* of the utilization in a resolution of one second, and hence the carrier wave f is not shown in the graph.

Algorithm 3 Signal $(bit,\ f,\ nCycles)$

```
1: if bit == 1 then
2:     for j ← 0 to nCycles do
3:         busywait((1000/f) * 0.5))
4:         sleep((1000/f) * 0.5)
5:     end for
6: else
7:     sleep(nCycles * (1000/f))
8: end if
```

Algorithm 4 Busywait (ms)

```
1: T1 ← getCurrentTime()
2: while (getCurrentTime() − T1 < ms) do ;
3: end while
```

transmission (Algorithm 1, lines 6-7). We used this technique in order to start the generation of the carrier wave at the same time in all worker threads. Each worker thread (*workerThread*) is bound to a different core (Algorithm 2, line 1), and waits for the beginning of the transmission (Algorithm 2, line 2). It then iterates on the stream of bits to transmit and invokes the signal generation function for each bit (Algorithm 2, line 4).

The *signal* function receives the bit to transmit (bit). In the case of logical '1' it repeatedly starts and stops the workload at the carrier frequency f for $nCycles$ cycles (Algorithm 3, lines 2-5). We overload the core using the busy waiting technique as presented in the *BusyWait* function. This function causes full utilization of the core for the time period and returns. In the case of logical '0' it sleeps for a period of $nCycles$ cycles (Algorithm 3, line 7).

Based on these algorithms, we implemented a transmitter for Linux OS. We used the *sched_setaffinity* system call to bind each thread to a specific CPU core [62]. For the threads synchronization we used the mutex object functions *pthread_mutex_...()* [63]. For thread sleeping we used the *sleep()* system call [64].

Figure 3 shows the CPU utilization of a transmitter running on a PC with eight cores. The four transmitting threads are bound to four cores (CPU1-CPU4) while the rest of cores remain available for other tasks.

Stealth: The transmitting code shown above requires no elevated privileges and can be initiated from an ordinary user space process. The code consists of basic CPU operations such as busy loops, which do not expose malicious behavior, making it highly evasive from static and dynamic (runtime) malware detection solutions.

TABLE I
AMPLITUDE-SHIFT KEYING

Symbol	Amplitude	# of cores
00	A_0	0 (or 1)
01	A_1	2
10	A_2	3
11	A_3	4

B. Data Modulation

By using different cycle times in the signal generation algorithm, we are able to control the carrier wave frequency. We also have some control of the carrier wave's amplitude by varying the number of cores used for generating the signal. Based on that, we implemented three different data modulation schemes for the transmission: Amplitude-shift keying (ASK), frequency-shift keying (FSK), and the more efficient orthogonal frequency-division multiplexing (OFDM) modulation. In the following sections, we describe each of the three modulations used.

1) Amplitude-Shift Keying: In amplitude-shift keying modulation the data is represented by the level of the amplitude of the carrier wave, whereas each level represents a different symbol. The transmitting code controls the strength (amplitude) of the magnetic field by using different numbers of cores for the transmissions. Accordingly, the number of cores available for the transmission is the number of symbols available. The relationship between the number of symbols available and the number of bits that can be represented by a symbol is $M = 2^n$ where M is the number of symbols, and n is the number of bits. Table I presents a case in which four CPU cores are available for the transmission. We encode the four symbols '00', '01', '10,' and '11' by four amplitude levels, A_0, A_1, A_2, and A_3, respectively. Figure 4 shows the waveform of a binary sequence ('11100100') modulated with four level ASK and transmitted from a desktop PC with four cores.

The On-Off Keying (OOK) modulation is the simplest form of ASK in which the data is represented by the presence/absence of the carrier wave. The presence of a carrier wave represents the symbol '1,' while its absence represents the symbol '0'.

2) Frequency-Shift Keying: In frequency-shift keying (FSK) the data is represented by a change in the frequency of a carrier

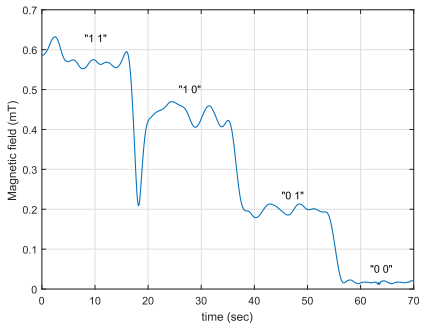

Fig. 4. The waveform of a binary sequence (‘111000100’) modulated with four amplitudes ASK.

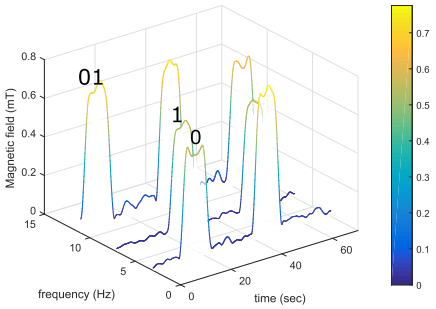

Fig. 5. The spectrogram of the binary sequence (‘010101010’) modulated with FSK via three frequencies (3 Hz, 7 Hz and 13 Hz).

wave. Recall that the transmitting code can determine the frequency of the generated magnetic signal. In FSK, each frequency represents a different symbol. Figure 5 shows the time-frequency spectrogram of a binary sequence (‘010101010’) modulated with three frequency FSK as transmitted from a PC with four cores. In this modulation, the frequencies 3Hz, 7Hz, and 13Hz have been used to encode the symbols ‘0’, ‘1,’ and ‘01’ respectively.

3) Orthogonal Frequency-Division Multiplexing: In orthogonal frequency-division multiplexing data is represented by multiple carrier frequencies in parallel. In our case, we use different cores to transmit data in different sub-carriers at a range of 0-50Hz. In each sub-carrier, we used OOK to modulate the data. Note that since the sub-carriers’ signals are generated in parallel, the maximal number of sub-carriers is equal to the number of cores available for the transmissions. Figure 6 presents the binary sequence (‘1101111011’) modulated with OFDM with two sub-carriers as transmitted from a PC with four cores. In this modulation, 7Hz (core 1) and 11Hz (core 2) have been used to encode the symbols ‘00’, ‘01’, ‘10’ and ‘11’.

C. Frames

We transmit the data in small packets composed of a preamble, a payload, and a parity bit/forward error correction (FEC) codes.

- **Preamble.** A preamble header is transmitted at the beginning of every packet. It consists of a sequence of four alternating bits (‘1010’) which helps the receiver determine the carrier wave frequency and amplitude. In addition, the preamble allows the receiver to detect the beginning of a transmission. Note that in our covert channel the amplitude of the carrier wave is unknown to the receiver in advance, and it mainly depends on what type of transmitting computer is used, the number of cores participating in the transmission, and the distance between the transmitter and the receiver. These parameters are synchronized with the receiver during the preamble.
- **Payload.** The payload is the raw data to be transmitted. In our case, we arbitrarily choose 32 bits as the payload size.
- **Parity bit.** For error detection, a parity bit is added to the end of the frame. The receiver calculates the parity for the received payload, and if it differs from the received parity bit, an error is detected. A more robust protocol may involve advanced error detection and error correction codes (e.g., cyclic redundancy checks). For simplicity we do not consider this in the current paper.
- **Forward error correction (FEC)** The magnetic covert channel may be operated in a noisy environment. Since the communication channel is unidirectional (e.g., no feedback or ACK messages) a forward error correction code might be added to the frame. We added FEC support in the transmitter and receiver. In our case we used the Reed-Solomon codes, which were applied to each frame. Our implementation is based on the RSCODE project which is an open-source library of a Reed-Solomon error correction algorithm [65].

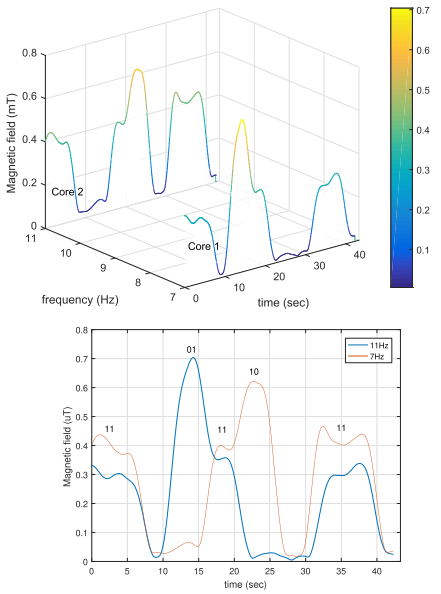

Fig. 6. The spectrogram (upper figure) and waveforms (lower figure) of the binary sequence (‘1101111011’) modulated with OFDM and two sub-carriers (7 Hz and 11 Hz).

VI. Analysis & Evaluation

In this section, we present an analysis and evaluation of the proposed covert channel. We evaluate the covert channel using various measures including:

- **Signal strength.** The strength of the generated signals given different number of threads.

TABLE II
THE COMPUTERS USED IN THE EXPERIMENTS

#	Name	Model	Motherboard/board	CPU	# of cores
1	PC-1 (desktop PC)	Infinity desktop PC	Gigabyte H87M- D3H	Intel Core i7-4770 CPU @ 3.4GHz	4 (8 logical)
2	PC-2 (desktop PC)	Lenovo desktop PC	Panda L-IQ45	Intel Core Quad-Q9550 CPU @ 2.83GHz	4 (4 logical)
3	Laptop	HP ProBook 650 G2	Intel	Intel Core i5-Q6200U CPU @ 2.4GHz	2 (4 logical)
4	Server	IBM System x3500 M4	Intel C602J	Intel Xeon CPU E5-2620	12 (24 logical)
5	NUK (small form factor)	Lenovo ThinkCentre M93p	Intel Q87 express ThinkCentre M93/M93p	Intel Core i7 -4785T	4 (8 logical)
6	IoT (IoT/embedded device)	Raspberry Pi 3	Raspberry Pi 3 model B V1.2	BCM2837 64-bit ARMv8, processor Cortex A53	4

- **Distances.** The strength of the received signals from various distances.
- Channel capacity. The theoretical bounds on the capacity of the communication channel for various setups.
- **Data transfer.** The actual bit rate and bit error rate (BER) of the received signals of various computers from various distances.
- **Virtual machines.** The interoperability of the covert channel from within a virtual machine (VM).
- **Interference.** The interference of the transmitting threads with various workloads of the CPU.

Our experiments focus on the evaluation of the CPU as an unintended, low power magnetic transmitter used for covert communication.

A. Experimental Setup

1) Transmitters (Computers): The experimental setup consists of six types of computers that are used for the transmissions: two off-the-shelf standard desktop PCs, a laptop computer, a small form factor computer, a server machine with multi-core processors, and a low power embedded device. Unless otherwise specified, the systems in the experiments were run with Linux Ubuntu 16.04 64-bit. A detailed list of the computers is provided in Table II.

2) Hyper-Threading: Note that modern Intel CPUs support Hyper-Threading Technology [66]. In this technology, each physical core exposes two logical (virtual) cores to the operating system. The CPU shares the workload between the logical cores when possible for better utilization. In the experiments, we bound the transmitting threads to the systems' logical cores rather than the physical cores, i.e., in a system with four physical cores and eight virtual cores we can potentially run eight concurrent transmitting threads.

3) Receiver: For the reception, we used the HMR2300 (Honeywell) magnetic sensor [67]. This is a digital magnetometer which is capable of sampling the magnetic field in three axes. It is in use in a wide range of applications such as compassing and navigation, traffic and vehicle detection, laboratory instrumentation, and security systems. The three internal magnetoresistive sensors are oriented in orthogonal directions to measure the X, Y, and Z vector components of a magnetic field, and the output is converted to 16-bit digital values using an internal analog-to-digital converter. The sensor resolution is approximately $70nT$ (nanotesla), and the sampling rate is up to 154 samples per second.

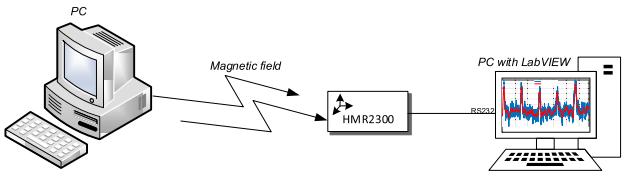

Fig. 7. The measurement setup.

4) Measurement Setup: The measurement setup is shown in Figure 7. The magnetic sensor is connected to the computer using a serial communication port (RS-232) which is configured to a full-duplex 19,200 data rate. The data is collected with a system driven by a LabVIEW data flow visual programming language.

5) Layouts: The layout of the transmitter and receiver during the experiments represents the attack scenarios described in Section III. The magnetic receiver is positioned near the transmitting computer, in the same room or located behind a wall. When the distance between the transmitter and receiver is a above 25cm, the magnetic receiver was put behind a wall.

B. Signal Strength

1) Number of Cores: As discussed, the number of cores used in the transmission directly influences the strength (amplitude) of the magnetic signal (i.e., more transmitting threads yield a stronger signal). Figure 8 shows the measurements of three different transmitters: a PC (PC-1), a laptop, and a server. In this test, we used one thread per core and set the carrier frequency to 20Hz. The magnetic field measured at a distance of 20cm from the transmitting computers showed a gradual increase when increasing numbers of threads are used. As expected, the twelve core server showed the greatest increase; from a magnetic field strength of 0.05mT (two threads) to a magnetic field strength of 0.9mT (twelve threads). The magnetic field of the PC increased from 0.15mT (one thread) to almost 0.6mT (eight threads). The magnetic field of the laptop showed almost no increase in the magnetic field strength between one and four threads.

2) Distance: The strength of a magnetic field decreases rapidly, inversely proportional to the third power of the distance from the magnetic source. Figure 9 and Table III show the strength of magnetic signals as measured at various distances from five transmitters. Note that Figure 9 shows the magnetic field in a logarithmic scale. Using the HMR3200 sensor, the magnetic signals were received at a maximal distance

TABLE III
MEASUREMENTS OF THE MAGNETIC FIELD OF FIVE TRANSMITTERS AT VARIOUS DISTANCES

	10 cm	20 cm	40 cm	60 cm	80 cm	100cm	120cm	140cm	150cm
PC-1	0.99 mT	0.51 mT	0.13 mT	0.042 mT	0.019 mT	0.013 mT	-	-	-
PC-2	0.63 mT	0.27 mT	0.059	0.022 mT	0.01 mT	0.009 mT	-	-	-
Laptop	0.084 mT	0.019 mT	0.012 mT	-	-	-	-	-	-
Server	0.6 mT	0.35 mT	mT 0.14	0.07 mT	0.035 mT	0.026 mT	0.019 mT	0.017 mT	0.013 mT
NUK	1.4 mT	0.8 mT	0.2 mT	0.016 mT	0.013 mT		-	-	-

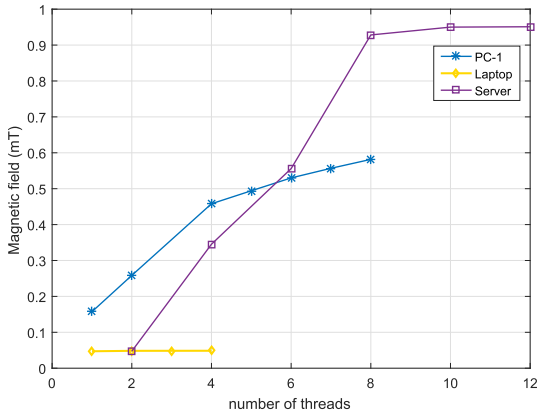

Fig. 8. Measurement of the magnetic signal generated by different numbers of threads on three computers: PC-1, laptop, and server.

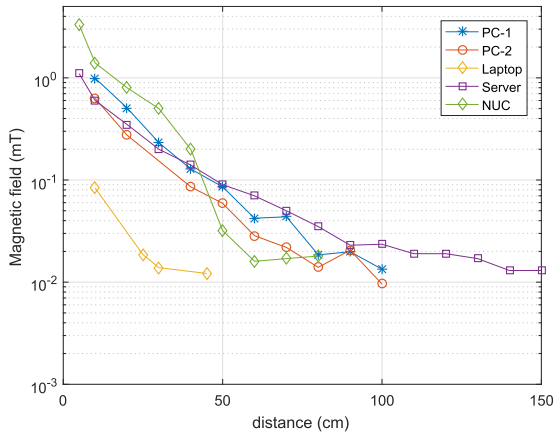

Fig. 9. A 20Hz signal generated by five computers (PC-1, PC-2, laptop, server, and NUK) measured at distances of 0 to 150 cm.

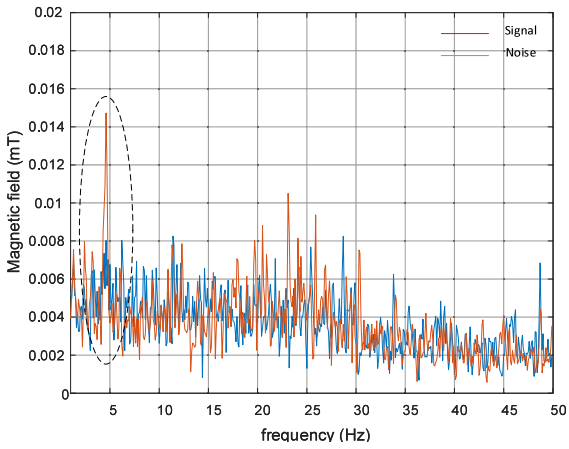

Fig. 10. A 4.7Hz signal transmitted from a Faraday shielded PC-1 as received from a distance of 100 cm away.

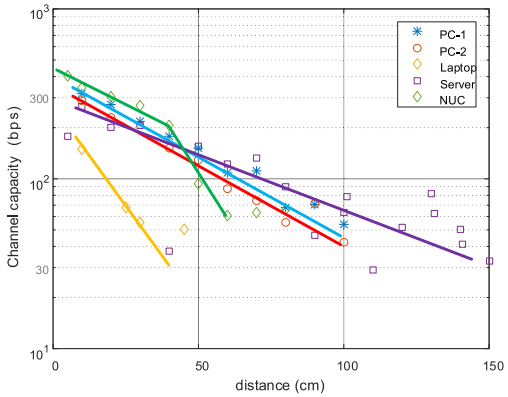

Fig. 11. The channel capacity of five transmitters based on the SNR measured at a range of distances.

of 50 cm for the laptop, 100 cm for the desktop PCs and small form factor PC, and 150 cm for the server. Given the resolution of the HMR2300 sensor and the signal to noise ratio (SNR), we stopped the measurements at a field power of 10^{-2} mT. Note that it is possible to increase the reception distance with more sensitive types of magnetic sensors (e.g., [68]). We left this direction of research for future work in this field.

In Section IV we showed that magnetic fields at very low frequencies have a low impedance and they can bypass metal shields. Figure 10 depicts this by showing a magnetic signal transmitted from within a Faraday cage. In this case, PC-1 was transmitting at 4.7Hz and located within a Faraday cage, 100 cm from the magnetic sensor. The blue line shows the background noise, while the red line shows the signal. As can be seen, the strength of the measured signal is 0.014 mT - significantly higher than a background noise of 0.007 mT in this band.

C. Channel Capacity

Using the Shannon-Hartley theorem, we calculated the maximum bitrate for this communication channel. Figure 11 shows the calculated channel capacity, given the quality of signals measured from PC-1, PC-2, laptop, server, and NUK. In our case, the bandwidth (B) is 50Hz given the low frequencies used for the transmissions. The signal (S) and noise (N) were calculated based on the SNR measurements taken for each of the computers at distances of 0 to 150 cm. As can be observed, for desktop, server and NUK computers, the channel

TABLE IV
BER MEASUREMENTS FOR PC-1, SERVER AND NUK

PC-1	**5 cm**	**20 cm**	**40 cm**	**60 cm**	**80 cm**	**100 cm**	**120 cm**
1 bit/sec	0%	0%	0%	0%	5%	10%	-
10 bit/sec	0%	0%	25%	-	-	-	-
40 bit/sec	0%	20%	-	-	-	-	-
Server	**0 cm**	**20 cm**	**40 cm**	**60 cm**	**80 cm**	**100 cm**	**120 cm**
1 bit/sec	0%	0%	0%	0%	0%	0%	18%
10 bit/sec	0%	0%	28%	-	-	-	-
40 bit/sec	0%	30%	-	-	-	-	-
NUC	**0 cm**	**20 cm**	**40 cm**	**60 cm**	**80 cm**	**100 cm**	**120 cm**
1 bit/sec	0%	0%	0%	0%	10%	20%	-
10 bit/sec	0%	0%	28%	-	-	-	-
40 bit/sec	0%	30%	-	-	-	-	-

capacity varies from 300 bit/sec to 30 bit/sec depending on the distance of the receiver. The channel capacity for the laptop is significantly lower due to the weak magnetic signals it generates.

D. Data Transfer

The channel capacity represents the upper theoretical limits of a communication channel. The actual bitrate is usually lower than the channel capacity and is determined by the modulation scheme and the quality of the transmitter and receiver used. We measured the bit error rate (BER) of PC-1, the server, and NUK computers for distances of 0 to 120 cm from the transmitting computer. In this test, all of the available cores were used for the data transmission. We tested the transmissions at three bitrates (1, 10, and 40 bit/sec) using the simple OOK modulation and stopped the tests when the results showed a BER of 30% or higher. For the BER measurements we transmitted randomly generated encryption keys in a size of 256-bit.

The results, presented in Table IV, show that up to a distance of 100 cm, the effective transmission rate is 1 bit/sec for the three computers, with a maximal BER of 10%. The higher transmission rates of 10 bit/sec and 40 bit/sec are feasible only when the sensor was in close proximity (5-20 cm away) to the transmitting computer. Note that it is possible to increase the distance by reducing the transmission rates further. However, for the evaluation we consider a transmission rate of 1 bit/sec as the minimal bitrate justifying this attack model. Figure 12 shows the waveforms of the transmitted key (‘10101110110101010110…’) encoded in OOK, as transmitted from the server computer during the BER measurements. The data was transmitted at a speed of 5 bit/sec and received at distances of 50 cm, 75 cm, and 150 cm away, with an SNR of 10dB, 8dB and 4.4dB respectively.

1) Embedded/IoT Device: Embedded devices usually consume just a small amount of power, hence emitting weak magnetic fields. Our measurements show that the proposed covert channel works with low power devices only when the magnetic sensor is in close proximity of the device. Figure 13 shows the waveform of an alternating binary sequence modulated with OOK, as transmitted from the Raspberry Pi 3. The data was transmitted at a speed of 41 bit/sec and received at distances of 10 cm away with a BER of 0% and a SNR of 15dB.

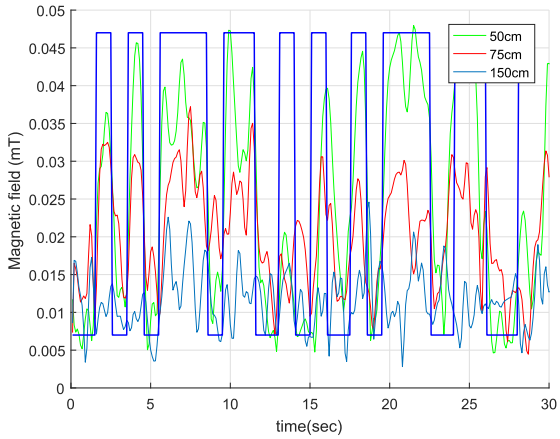

Fig. 12. Server BER measurements for distances of 50, 75, and 150 cm (with BER of 0%, 0%, and 22%, respectively).

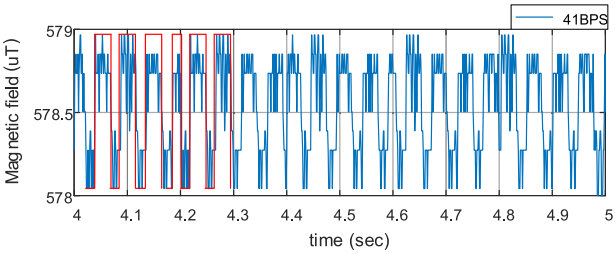

Fig. 13. The waveform of an alternating binary sequence modulated with OOK, as transmitted from the Raspberry Pi 3 at 41 bit/sec.

E. Virtual Machines

Virtualization technologies are widely used in modern IT environments, including in desktop/server virtualization systems and for private and public clouds. One of the advantages of virtualization is the resource isolation it provides. Virtual Machine Monitors (VMM) and hypervisors provide a separation between the guest operating system and hardware resources. We examined the operability of a transmitter running in a virtualized environment. Our main goal was to determine whether the execution of the transmitted threads on virtualized CPU cores caused interruptions or delays which may affect the signal generation. Figure 14 shows the waveforms of two signals transmitted from PC-1. The first signal was generated from the host computer, and the second signal was generated from within a virtual machine. The receiver was located 30cm from the transmitting computer. Both signals depict the transmission of a random sequence of bits for a duration of 30 seconds. Both the guest and the host were running Linux Ubuntu 16.04 64-bit. We used VMWare Workstation Player 14.0 for the virtualization and configured the host machine to support four CPU processors. As can be seen, the magnetic signal generated from the VM is highly correlated to the magnetic signal generated directly from the host computer, both having an SNR of 15dB. We repeated this test 20 times every time with a random sequence of bits and received the same results. More specifically, we experienced no time delay or reduction in the power of the signal when it was generated from the VM.

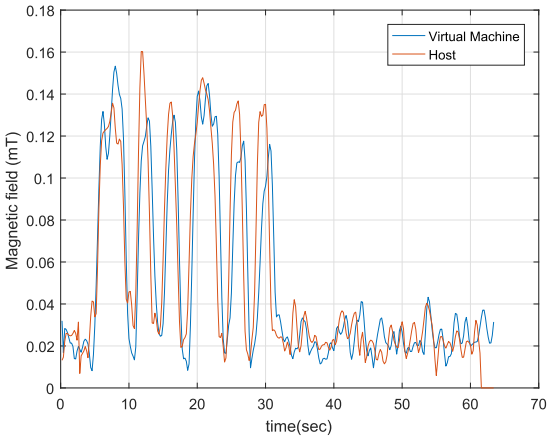

Fig. 14. The waveforms of two signals transmitted from PC-1 (VM/Host).

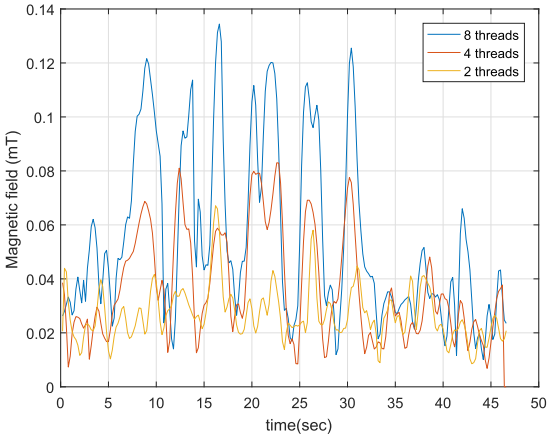

Fig. 15. The waveform of a signal generated from a virtual machine using 2, 4, and 8 threads.

We also investigated the feasibility of controlling the signal strength from virtual machines by using different numbers of cores. Figure 15 presents the waveform of a signal generated from a VM with the same setup as described above, using 2, 4, and 8 threads. As can be seen, the transmissions of 2, 4, and 8 threads generate signals at 0.12mT, 0.08mT, and 0.04mT, with SNR levels of 15dB, 11dB and 5dB respectively. Generally, employing different numbers of virtual cores in a VM yields different levels of signals, similar to host-based transmissions. In the context of the communication channel, it allows the attacker to use amplitude modulations from VMs.

F. Interference With Processes

The threads that generate the magnetic signals share the CPU time with other processes in the operating system. We examined whether the activities of various processes in the system interfere with the signal generation. For this evaluation, we run the transmitting process in PC-1, while employing the following five types of workloads commonly run on desktop PCs:

1) The system was idle and only the default processes were running in the background.

TABLE V
INTERFERENCE WITH VARIOUS WORKLOADS

#	Workload	Application/Process	SNR (dB)
#1	Idle	Background processes	36.47
#2	Word processing	LibreOffice Writer	36.00
#3	Video playing	VLC Player	35.31
#4	Backup	rsync	32.04
#5	Calculations	matho-primes	33.97

2) The *LibreOffice Writer* [69] was open, and the user typed a document.
3) The *VLC media player* [70] was playing an HD video clip.
4) The Linux *rsync* command was performing a backup of local folders in the HDD [71].
5) Linux *matho-primes* program was performing the calculations of big prime numbers [72].

Table V summarizes the SNR measured for each of the five workloads. We used eight threads for the transmission of an alternating bit sequence ('10101010') using OOK modulation. The receiver was located 35cm from the transmitting computer. The signals depict the transmission of the alternating bit sequence for a duration of 30 seconds. Naturally, the idle state where no other processes interfered with the transmitting process yielded the strongest signal with a measured SNR of 36dB. The word processing and video playing processes consumed just small slices of the CPU time and hence reduced the signal strength at an intermediate level with an SNR of 35dB to 36dB. The calculation and backup workloads caused the greatest degradation in the received signals due to the intensive CPU and I/O operations they perform. The SNR in these cases was reduced to levels of 32 to 34dB.

The results show that the proposed covert channel is usable even when other active processes are running in the system. However, CPU intensive operations add noise to the generated signal, hence decreasing the effective range and increasing the bit error rate of the transmissions.

VII. COUNTERMEASURES

A. Detection

Detection of covert channels could take place by security systems running on the computer. In this approach, security solutions such as AVs that continuously trace the activities of a computer's processes and try to detect malicious operations; in the case of a magnetic covert channel, a process that abnormally regulates the CPU workload would be tagged as suspicious. However, many types of applications use working threads that affect the processor's workload, and therefore, such a detection approach would likely suffer from a high rate of false alarms. As mentioned in Section V, the signal generation involves simple CPU operations basic system calls. Tracing such non-privileged CPU instructions at runtime necessitates executing the processes in a single step mode, which can severely degrade system performance [20]. Another approach is to detect the covert channel externally, by monitoring the magnetic field in the area of the computer. The magnetic field is measured and analyzed to find deviations

TABLE VI
LIST OF TECHNOLOGICAL COUNTERMEASURES

Countermeasure	Description	Drawbacks/challenges
Malicious activity detection (software)	Detecting the transmitting threads	False positives; can be bypassed (e.g., by rootkits)
Magnetic activity detection (hardware)	Detecting abnormal magnetic fields	False positives; expensive
Ferromagnetic shielding (mu-metal)	Shielding with Ferromagnetic material	Expensive
Magnetic field jammer (hardware)	Jamming (magnetic field generator)	Expensive
Field cancellation	Producing the counter magnetic fields	Expensive
Random workload generator	Interrupting the transmissions	Degrades system performance; can be bypassed

from the standards. Note that this approach requires a magnetic sensor and additional hardware in the proximity of each monitored computer.

B. Prevention

There are three different approaches that can be used to prevent attackers from establishing a magnetic covert channel: shielding, jamming, and zoning.

1) Shielding: Shielding computers, effectively enclosing them to protect them from low frequency magnetic fields, is considered impractical except for special military or scientific purposes. As discussed in the evaluation section, magnetic fields lower than 50Hz have very low impedance and are difficult to reduce, since this would require very thick metal shielding. For effective magnetic shielding, Ferromagnetic materials such as mu-metal should be used [73]. Ferromagnetic materials require less thick shielding, and hence are more practical for the construction of shielded computer enclosures; however, it is difficult to provide effective magnetic shielding against low frequencies even with Ferromagnetic material [59]. Magnetically shielded rooms provide shielding protection on a larger scale. These rooms, which consist of several layers of Ferromagnetic plates, are expensive and weigh several tons. For a more in depth discussion of different approaches for magnetic shielding, we refer the interested reader to [74].

2) Signal Jamming: Signal jamming is commonly used to mitigate electromagnetic and acoustic covert channels [75]. In this approach, a strong signal that interferes with unauthorized communication is generated in the area requiring protection. The same approach can be used for magnetic communication. Commercial magnetic field generators such as MGA 1030 can generate magnetic fields as strong as 1000 A/m at low frequencies (up to 1kHz) [76]. The power of such a magnetic field is hundreds of times stronger than the magnetic field generated by the CPU, and therefore overrides its magnetic signals. Field cancellation, also known as active magnetic shielding, is another type of signal jamming which is unique to magnetic emanation. This technique uses special equipment that monitors magnetic fields and cancels them by driving a current that produces counter magnetic fields [77]. An interesting software level jamming solution is to execute background processes that initiate random magnetic transmissions. The random signals interfere with the transmissions of the malicious process, however random workloads weaken system performance and may be infeasible in some environments (e.g., real-time systems).

3) Zoning: Procedural countermeasures involve a physical separation of emanating equipment from potential receivers. This approach is referred to as 'zoning' in the National Security Telecommunications and Information Systems Security Advisory Memoranda (NSTISSAM) and NATO standards. For example, the SDIP-27 and SDIP-28 NATO standards define separated zones in which electronic equipment is allowed [78]. In these standards, sensitive computers are kept in restricted areas in which certain equipment is banned. In our case, magnetic receivers of any kind should be banned in the proximity of the sensitive computers.

The detection and prevention based countermeasures and their limitations/weaknesses are summarized in Table VI.

VIII. CONCLUSION

This paper presents a new type of covert channel based on low frequency magnetic fields. This method allows attackers to exfiltrate data from isolated, air-gapped computers to a nearby magnetic sensor. Moreover, due to the nature of low frequency magnetic fields, they penetrate through metals. This makes our covert channel possible even in a constrained environment where the computers are enclosed within Faraday shielding and the conventional electromagnetic covert channels fail. We present scientific background and explain the characteristics of magnetic fields and the signal generation technique. We introduce a malware codenamed 'ODINI,' which controls the magnetic fields emitted from the computer by controlling the workload of the CPU cores. We show that the malware can work from a user-level process and can operate from within an isolated virtual machine (VM), without requiring special execution privileges. We evaluate the covert channel and show that it works on a wide range of computers. We also propose different types of defensive countermeasures to detect and prevent this threat. Our results show that data can be successfully exfiltrated from air-gapped, Faraday caged systems via low frequency magnetic fields at bitrates of 1-40 bit/sec.

ACKNOWLEDGMENT

Demonstration video: https://www.youtube.com/watch?v=h07iXD-aSCA.

REFERENCES

[1] E. MacAskill, S. Thielman, and P. Oltermann. (2017). WikiLeaks Publishes 'Biggest Ever Leak of Secret CIA Documents.' The Guardian. [Online]. Available: https://www.theguardian.com/media/2017/mar/07/wikileaks-publishes-biggest-ever-leakof-secret-cia-documents-hacking-surveillance

[2] K. Zetter, "Sony got hacked hard: What we know and don't know so far," *Wired*, Jan. 2014.

[3] S. Thielman. (Dec. 2016). Yahoo Hack: 1 bn Accounts Compromised by Biggest Data Breach in History. The Guardian. [Online]. Available: https://www.theguardian.com/technology/2016/dec/14/yahoo-hack-security-of-one-billion-accounts-breached

[4] M. Guri, M. Monitz, and Y. Elovici, "Bridging the air gap between isolated networks and mobile phones in a practical cyber-attack," *ACM Trans. Intell. Syst. Technol.*, vol. 8, no. 4, p. 50, 2017.

[5] E. Byres, "The air gap: Scada's enduring security myth," *Commun. ACM*, vol. 56, no. 8, pp. 29–31, 2013.

[6] Wikipedia. *Classified United States Website*. Accessed: May 4, 2018. [Online]. Available: https://en.wikipedia.org/wiki/Classified_United_States_website

[7] M. Maybury *et al.*, "Analysis and detection of malicious insiders," MITRE, Bedford, MA, USA, Tech. Rep., 2005.

[8] S. Abraham and I. Chengalur-Smith, "An overview of social engineering malware: Trends, tactics, and implications," *Technol. Soc.*, vol. 32, no. 3, pp. 183–196, 2010.

[9] *ProjectSauron: Top Level Cyber-Espionage Platform Covertly Extracts Encrypted Government Comms*. Accessed: Apr. 5, 2018. [Online]. Available: https://securelist.com/faq-the-projectsauron-apt/75533/

[10] *A Fanny Equation: 'I am Your Father, Stuxnet'*. Accessed: May 4, 2018. [Online]. Available: https://securelist.com/a-fanny-equation-i-am-your-father-stuxnet/68787/

[11] *The Regin Platform Nation-State Ownage of GSM Networks*. Accessed: Apr. 5, 2018. [Online]. Available: https://securelist.com/files/2014/11/Kaspersky_Lab_whitepaper_Regin_platform_eng.pdf

[12] R. Langner, "Stuxnet: Dissecting a cyberwarfare weapon," *IEEE Security Privacy*, vol. 9, no. 3, pp. 49–51, May/Jun. 2011.

[13] R. Grant, "The cyber menace," *Air Force Mag.*, vol. 92, no. 3, pp. 1–6, 2009.

[14] *Hammerdrill v2.0*. Accessed: Apr. 5, 2018. [Online]. Available: https://wikileaks.org/ciav7p1/cms/page_17072172.html

[15] Wikileaks. *Vault 7: Projects*. Accessed: Apr. 5, 2018. [Online]. Available: https://wikileaks.org/vault7/?#Brutal%20Kangaroo

[16] M. Guri, G. Kedma, A. Kachlon, and Y. Elovici, "AirHopper: Bridging the air-gap between isolated networks and mobile phones using radio frequencies," in *Proc. 9th Int. Conf. Malicious Unwanted Softw., Americas (MALWARE)*, Oct. 2014, pp. 58–67.

[17] M. G. Kuhn and R. J. Anderson, "Soft tempest: Hidden data transmission using electromagnetic emanations," in *Information Hiding*, vol. 1525. Cambridge, U.K.: Univ. Cambridge, 1998, pp. 124–142.

[18] M. G. Kuhn, "Compromising emanations: Eavesdropping risks of computer displays," Ph.D. dissertation, Univ. Cambridge, Berkeley, Berkeley, CA, USA, 2002.

[19] M. Vuagnoux and S. Pasini, "Compromising electromagnetic emanations of wired and wireless keyboards," in *Proc. USENIX Secur. Symp.*, 2009, pp. 1–16.

[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 *Proc. USENIX Secur. Symp.*, 2015, pp. 849–864.

[21] M. Guri and Y. Elovici, "Bridgeware: The air-gap malware," *Commun. ACM*, vol. 61, no. 4, pp. 74–82, Mar. 2018. [Online]. Available: http://doi.acm.org/10.1145/3177230

[22] L. O. Hoeft and J. S. Hofstra, "Measured electromagnetic shielding performance of commonly used cables and connectors," *IEEE Trans. Electromagn. Compat.*, vol. 30, no. 3, pp. 260–275, Aug. 1988.

[23] Emp. *Emi Shielded Racks*. Accessed: May 4, 2018. [Online]. Available: https://hollandshielding.com/EMP-EMI-shielded-racks

[24] *EMF Shielding Devices for Computers and TVs*. Accessed: May 4, 2018. [Online]. Available: https://www.lessemf.com/computer.html

[25] *RF Shielded Rooms*. Accessed: May 4, 2018. [Online]. Available: http://www.comtest.eu/products/rf-shielded-rooms-doors/rf-shielded-rooms.html

[26] W. Mazurczyk, S. Wendzel, S. Zander, A. Houmansadr, and K. Szczypiorski, *Information Hiding in Communication Networks: Fundamentals, Mechanisms, Applications, and Countermeasures*, vol. 7. Hoboken, NJ, USA: Wiley, 2016.

[27] K. Cabaj, L. Caviglione, W. Mazurczyk, S. Wendzel, A. Woodward, and S. Zander, "The new threats of information hiding: The road ahead," *IT Prof.*, vol. 20, no. 3, pp. 31–39, 2018.

[28] W. Mazurczyk and L. Caviglione, "Information hiding as a challenge for malware detection," 2015, *arXiv:1504.04867*. [Online]. Available: https://arxiv.org/abs/1504.04867

[29] S. J. Murdoch and S. Lewis, "Embedding covert channels into TCP/IP," in *Information Hiding*, vol. 3727. Springer, 2005, pp. 247–261.

[30] S. Zander, G. Armitage, and P. Branch, "A survey of covert channels and countermeasures in computer network Protocols," *IEEE Commun. Surveys Tuts.*, vol. 9, no. 3, pp. 44–57, 3rd Quart., 2007.

[31] B. Carrara and C. Adams, "Out-of-band covert channels—A survey," *ACM Comput. Surv.*, vol. 49, no. 2, p. 23, 2016.

[32] *Tempest for Eliza*. Accessed: May 4, 2018. [Online]. Available: http://www.erikyyy.de/tempest/

[33] M. Guri, M. Monitz, and Y. Elovici, "USBee: Air-gap covert-channel via electromagnetic emission from USB," in *Proc. 14th Annu. Conf. Privacy, Secur. Trust (PST)*, Dec. 2016, pp. 264–268.

[34] (2015). *Funtenna*. Accessed: May 4, 2018. [Online]. Available: https://github.com/funtenna

[35] A. Madhavapeddy, R. Sharp, D. Scott, and A. Tse, "Audio networking: The forgotten wireless technology," *IEEE Pervasive Comput.*, vol. 4, no. 3, pp. 55–60, Jul. 2005.

[36] M. Hanspach and M. Goetz, "On covert acoustical mesh networks in air," 2014, *arXiv:1406.1213*. [Online]. Available: https://arxiv.org/abs/1406.1213

[37] B. Carrara and C. Adams, "On acoustic covert channels between air-gapped systems," in *Proc. Int. Symp. Found. Pract. Secur.* Springer, 2014, pp. 3–16.

[38] M. Guri, Y. Solewicz, A. Daidakulov, and Y. Elovici, "Fansmitter: Acoustic data exfiltration from (speakerless) air-gapped computers," 2016, *arXiv:1606.05915*. [Online]. Available: https://arxiv.org/abs/1606.05915

[39] M. Guri, Y. Solewicz, A. Daidakulov, and Y. Elovici, "Acoustic data exfiltration from speakerless air-gapped computers via covert hard-drive noise (diskfiltration)," in *Proc. Eur. Symp. Res. Comput. Secur.* Oslo, Norway: Springer, 2017, pp. 98–115.

[40] J. Loughry and D. A. Umphress, "Information leakage from optical emanations," *ACM Trans. Inf. Syst. Secur.*, vol. 5, no. 3, pp. 262–289, 2002.

[41] 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*. Cham, Switzerland: Springer, 2017, pp. 161–184. doi: 10.1007/978-3-319-60876-1_8.

[42] M. Guri, B. Zadov, A. Daidakulov, and Y. Elovici, "xLED: Covert data exfiltration from air-gapped networks via switch and router LEDs," in *Proc. 16th Annu. Conf. Privacy, Secur. Trust (PST)*, 2018, pp. 1–12.

[43] M. Guri, O. Hasson, G. Kedma, and Y. Elovici, "An optical covert-channel to leak data through an air-gap," in *Proc. 14th Annu. Conf. Privacy, Secur. Trust (PST)*, 2016, pp. 642–649.

[44] A. C. Lopes and D. F. Aranha, "Platform-agnostic low-intrusion optical data exfiltration," in *Proc. ICISSP*, 2017, pp. 474–480.

[45] M. Guri and D. Bykhovsky, "Air-jumper: Covert air-gap exfiltration/infiltration via security cameras & infrared (ir)," *Comput. Secur.*, vol. 82, pp. 15–29, May 2019.

[46] M. Guri, M. Monitz, Y. Mirski, and Y. Elovici, "BitWhisper: Covert signaling channel between air-gapped computers using thermal manipulations," in *Proc. IEEE 28th Comput. Secur. Found. Symp. (CSF)*, Jul. 2015, pp. 276–289.

[47] J. J. Sojdehei, P. N. Wrathall, and D. F. Dinn, "Magneto-inductive (MI) communications," in *Proc. MTS/IEEE Conf. Exhibit.*, vol. 1, Nov. 2001, pp. 513–519.

[48] *Through-the-Earth Two-Way Emergency Wireless Communications for Mine Industry Safety*. Accessed: May 4, 2018. [Online]. Available: http://www.teslasociety.ch/info/magnetlink/2.pdf

[49] R. Bansal, "Near-field magnetic communication," *IEEE Antennas Propag. Mag.*, vol. 46, no. 2, pp. 114–115, Apr. 2004.

[50] N. Matyunin, J. Szefer, S. Biedermann, and S. Katzenbeisser, "Covert channels using mobile device's magnetic field sensors," in *Proc. 21st Asia South Pacific Design Autom. Conf. (ASP-DAC)*, Jan. 2016, pp. 525–532.

[51] Kaspersky Lab. *Industrial Defence in-Depth*. Accessed: May 4, 2018. [Online]. Available: https://www.sans.org/summit-archives/file/summit-archive-1493412875.pdf

[52] *The Epic Turla (Snake/Uroburos) Attacks|Virus Definition|Kaspersky Lab*. Accessed: May 4, 2018. [Online]. Available: https://www.kaspersky.com/resource-center/threats/epic-turla-snake-malware-attacks

[53] GReAT, "Red october diplomatic cyber attacks investigation," Kaspersky SecureList, Tech. Rep., Jan. 2013.

[54] J. Rutkowska and A. Tereshkin, "Evil maid goes after truecrypt," *Invisible Things Lab*, 2009.

[55] V. P. Kodali, *Engineering Electromagnetic Compatibility: Principles, Measurements, Technologies, and Computer Models*. Piscataway, NJ, USA: Institute of Electrical and Electronics Engineers, 2001.

[56] M. N. O. Sadiku, *Elements of Electromagnetics*. London, U.K.: Oxford Univ. Press, 2014.
[57] V. Bindar, M. Popescu, and R. Craciunescu, "Aspects of electromagnetic compatibility as a support for communication security based on TEMPEST evaluation," in *Proc. 10th Int. Conf. Commun. (COMM)*, May 2014, pp. 1–4.
[58] C. L. Holloway, D. A. Hill, R. A. Dalke, and G. A. Hufford, "Radio wave propagation characteristics in lossy circular waveguides such as tunnels, mine shafts, and boreholes," *IEEE Trans. Antennas Propag.*, vol. 48, no. 9, pp. 1354–1366, Sep. 2000.
[59] *EMC for Systems and Installations Part 4—Filtering and Shielding*. Accessed: May 4, 2018. [Online]. Available: http://www.compliance-club.com/archive/keitharmstrong/systems_installations4.html
[60] V. Kelha, J. Pukki, R. Peltonen, A. Penttinen, R. Ilmoniemi, and J. Heino, "Design, construction, and performance of a large-volume magnetic shield," *IEEE Trans. Magn.*, vol. MAG-18, no. 1, pp. 260–270, Jan. 1982.
[61] J. von Kistowski, H. Block, J. Beckett, C. Spradling, K.-D. Lange, and S. Kounev, "Variations in CPU power consumption," in *Proc. 7th ACM/SPEC Int. Conf. Perform. Eng.*, 2016, pp. 147–158.
[62] *Sched_Setaffinity(2)—Linux Man Page*. Accessed: May 4, 2018. [Online]. Available: https://linux.die.net/man/2/sched_setaffinity
[63] *Pthread_Mutex_Lock(3): Lock/Unlock Mutex—Linux Man Page*. Accessed: May 4, 2018. [Online]. Available: https://linux.die.net/man/3/pthread_mutex_lock
[64] *Sleep(3)—Linux Manual Page*. Accessed: May 4, 2018. [Online]. Available: http://man7.org/linux/man-pages/man3/sleep.3.html
[65] *Rscode Project*. Accessed: Jul. 2, 2019. [Online]. Available: http://rscode.sourceforge.net/
[66] D. Marr *et al.*, "Hyper-threading technology in the netburst R microarchitecture," in *Proc. 14th Hot Chips*, 2002, pp. 1–37.
[67] *Honeywell HMR2300*. Accessed: May 4, 2018. [Online]. Available: https://aerocontent.honeywell.com/aero/common/documents/myaerospacecatalog-documents/Missiles-Munitions/HMR2300.pdf
[68] H. Auster *et al.*, "The THEMIS fluxgate magnetometer," in *The THEMIS Mission*. Springer, 2009, pp. 235–264.
[69] *Home | LibreOffice—Free Office Suite—Fun Project—Fantastic People*. Accessed: Jan. 11, 2018. [Online]. Available: https://www.libreoffice.org/
[70] VideoLAN. *Official Download of VLC Media Player, the Best Open Source Player*. Accessed: Jan. 11, 2018. [Online]. Available: https://www.videolan.org/vlc/index.html
[71] *Rsync(1)—Linux Man Page*. Accessed: Jan. 14, 2018. [Online]. Available: https://linux.die.net/man/1/rsync
[72] *Ubuntu Manpage: Matho-Primes—Generate Consecutive Prime Numbers*. Accessed: Jan. 11, 2018. [Online]. Available: http://manpages.ubuntu.com/manpages/zesty/man1/matho-primes.1.html
[73] H. J. ter Brake, H. Wieringa, and H. Rogalla, "Improvement of the performance of a mu-metal magnetically shielded room by means of active compensation (biomagnetic applications)," *Meas. Sci. Technol.*, vol. 2, no. 7, p. 596, 1991.
[74] V. V. Yashchuk, S.-K. Lee, and E. Paperno, *Magnetic Shielding*. Cambridge, U.K.: Cambridge Univ. Press, 2013, pp. 225–248.
[75] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "RFReact—A real-time capable and channel-aware jamming platform," *ACM SIGMOBILE Mobile Comput. Commun. Rev.*, vol. 15, no. 3, pp. 41–42, 2011.
[76] *Schloeder-EMV: Magnetic Field Generator and Analyzer*. Accessed: May 4, 2018. [Online]. Available: http://www.schloeder-emv.de/en/emc-products/emc-test-and-measurement-system/emc-generators-measurement-systems/magnetic-field-generator-and-analyzer.html
[77] Y. Okazaki, S. Yanase, and N. Sugimoto, "Active magnetic shielding with magneto-impedance sensor," *Int. J. Appl. Electromagn. Mech.*, vol. 13, nos. 1–4, pp. 437–440, 2001.
[78] R. J. Anderson, "Emission security," in *Security Engineering*. Hoboken, NJ, USA: Wiley, 2008, pp. 523–546.

Mordechai Guri received the B.Sc. and M.Sc. degrees in computer science from The Hebrew University of Jerusalem and the Ph.D. degree from the Department of Information Systems Engineering, Ben-Gurion University of the Negev (BGU), under the supervision of Prof. Y. Elovici. He is currently the Head of Research and Development at the Cyber-Security Research Center, BGU. He manages a research team in various topics of cyber security. His research interests include advanced malware, rootkits, air-gap security, covert channels, embedded systems, and mobile security.

Boris Zadov is currently a Security Researcher with the Cyber-Security Research Center, Ben-Gurion University of the Negev. He is also an electrical engineer and an expert in the domain of magnetic fields.

Yuval Elovici received the B.Sc. and M.Sc. degrees in computer and electrical engineering from the Ben-Gurion University of the Negev (BGU) and the Ph.D. degree in information systems from Tel Aviv University. He is currently the Director of the Telekom Innovation Laboratories, BGU, the Head of the Cyber Security Research Center, BGU, the Research Director of iTrust, SUTD, and a Professor with the Department of Information Systems Engineering, BGU. He has published articles in leading peer-reviewed journals and in various peer-reviewed conferences. He has also coauthored a book on social network security and a book on information leakage detection and prevention. His primary research interests are computer and network security, cyber security, web intelligence, information warfare, social network analysis, and machine learning.