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

Mordechai Guri, Boris Zadov, Andrey Daidakulov, Yuval Elovici
Ben-Gurion University of the Negev
Cyber Security Research Center
gurim@post.bgu.ac.il; zadov@post.bgu.ac.il; daidakul@post.bgu.ac.il; elovici@bgu.ac.il

Video: https://cyber.bgu.ac.il/advanced-cyber/airgap
(gurim@post.bgu.ac.il)

ABSTRACT

Air-gapped computers are computers which 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 by an eavesdropping adversary remotely. 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 an exploitation of the magnetic field generated by the computer's CPU. Unlike electromagnetic radiation (EMR), low frequency magnetic radiation propagates though the air, penetrating metal shielding such as Faraday cages (e.g., compass still works inside Faraday cages). Since the CPU is an essential part of any computer, our 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 code-named '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 receiver ('bug') placed nearby. We provide technical background and examine the characteristics of the magnetic fields. 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. We present signal generation, transmission algorithms, and discuss data encoding and modulation schemes. We analyze the covert channel and evaluate it on various types of computers. Finally, we propose different types of defensive countermeasures such as signal detection and signal jamming to cope with this type of threat.

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 in 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 in connected to the Internet, a motivated adversary will find a way to breach the network, evade security mechanisms, access sensitive data, and send it out to the attacker side (e.g., Vault 7 [1] , Sony [2], and Yahoo [3]).

A. Air-Gap Isolation

When highly sensitive data is involved, an organization may resort to so-called 'air-gap' isolation. In this approach, any type of physical or logical connection between the local network and the Internet is banned. The air-gap separation is maintained by enforcing strict regulations such as forbidding connectivity to unauthorized equipment and hardening the workstations in

the network. Today, air-gapped networks are used in military and defense systems, critical infrastructure, the finance sector, and other industries [4], [5]. Two examples of air-gapped networks are the NSANet and the Joint Worldwide Intelligence Communications System (JWICS), classified networks belonging to the United States' Defense Intelligence Agency [6]. However, even air-gapped networks are not immune to breaches. In the past decade, it has been shown that attackers can successfully penetrate air-gapped networks by using complex attack vectors, such as supply chain attacks, malicious insiders, and social engineering [7]–[9]. For example, in 2017 WikiLeaks published a reference to an hacking tool dubbed 'Brutal Kangaroo,' used to infiltrate air-gapped computers via USB drives [10]. This tool was used by the attackers to infect Internet workstations of the employees in the organization, and then wait for an employee to insert the USB drive into an air-gapped computer. Using these tools, attackers can breach the network, bypassing security systems such as AVs, firewalls, intrusion detection and prevention systems (IDS/IPS), and the like.

B. Air-Gap Covert Channels

After deploying a malware in the air-gapped network, the attacker may, at some point, wish to retrieve information such as encryption keys, documents, or passwords a behavior commonly used in APTs. 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. For example, electromagnetic covert channels have been studied for at least twenty years. In this type of communication, a malware modulates binary information over the electromagnetic waves radiating from computer components: LCD screens, communication cables, computer buses, and hardware peripherals [11]–[15].

C. Faraday Shielding

To cope with this type of leakage, particularly 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 and interferences [16]. Faraday shields may be deployed in a size of small enclosures to protect entire systems such as desktop PCs and display screens [17], [18], but they may also be used to protect entire rooms and even buildings [19]. 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 computers 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

- **Air-gap covert channel.** 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.
- **Leaking through Faraday shielding.** We introduce a covert channel that can evade the Faraday isolation. That is, it can work in highly secured systems which are kept within a Faraday cages where other types of covert channels (e.g., electromagnetic) fail. As far as we know, this is the first work that discusses the topic of Faraday cages and their evasion concerning covert channels.
- **Bypassing virtual machine (VM) isolation.** Virtual machines are often used as a security measure to add a layer of network isolation between the VM and the external environment. We show that the covert channel works even when the malicious code is executed in an isolated VM.
- **Discussion 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.

Our contribution concerning the characteristics of the malware is as follows:

- **Hardware availability.** The covert channel uses the standard CPU cores to generate and control the magnetic emanation. This makes the covert channel available on virtually any computer and device with a CPU.

- **Stealth.** The transmitting code is considered highly evasive, since it does not perform special CPU instructions or invoke specified API calls. This makes it difficult for AVs and anomaly detection systems to identify the malicious behavior of the malicious transmitter.
- **Required privileges.** The malware requires no special privileges (e.g., root or admin), and any user-level process can execute it in the system.

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

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. 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 [20]–[22]. In cases where there is no direct connection with the target network, the attacker may resort to so-called air-gap covert channels. These covert channels can be classified into five main categories: electromagnetic, magnetic, acoustic, thermal, and optical.

A. Electromagentic

The electromagnetic based covert channel has been the most researched topic in this field for at least twenty years. In 1988, Kuhn and Anderson [12] showed how attackers can exfiltrate data over the air-gap by controlling the electromagnetic waves emanating from display screens. In their method, called 'soft tempest,' a malicious code encodes information over AM signals generated by certain dither pixel patterns. Based on this work, in 2001 Thiele [23] 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 [4], [11], malware that can leak data from air-gapped networks to nearby mobile phones using controllable electromagnetic signals in the FM radio band emanating

from the video cable. Later on, in 2015, Guri et al presented GSMem [15], malware that leaks data from air-gapped computers using frequencies in the cellular band emitting from memory buses. 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 [24], the data bus, and GPIO ports [25] to generate covert electromagnetic signals for data exfiltration.

B. Acoustic, Optical, and Thermal

Acoustic. Exfiltrating data from air-gapped computers via acoustic signals has also been proposed. Hanspach discussed using ultrasonic sound waves (18-22kHz) to transmit data between air-gapped laptops using their speakers and microphones [26], however these methods are not relevant when speakers or microphones are not present. In 2016, Guri et al presented Fansmitter [27] and DiskFiltration [28], 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.

Optical. Data exfiltration using optical signals is another type of covert channel. In 2002, 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 [29]. More recently, Guri et al presented a covert channel that uses the hard drive indicator LED [30] the router LEDs [31] in order to leak data from air-gapped computers and networks. VisiSploit [32] 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 et al [33] presented a covert channel based on signals transmitted from IR LEDs in external USB devices attached to the computer. In 2017, Guri et al presented a method that uses the IR LEDs present in surveillance and security cameras to exfiltrate and infiltrate air-gapped networks remotely [34].

Thermal. In 2015, Guri et al presented a thermal based method called BitWhisper [35]. In this technique, an attacker can established 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.

C. Magnetic

Magnetic communication by itself is a known topic of research [36]. For example, the MagneLink Magnetic

Communication System (MCS) is a system which provides through-the-earth emergency wireless communication based on magnetic fields [37] Near-field magnetic induction (NFMI) communication is another type of magnetic method that allows short range communication between devices [38]. 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, Myhayun suggested using the hard disk drives (HDD) magnetic head to generate magnetic emissions, which can be sensed by a nearby smartphones magnetic sensor [39]. The smartphone needs to be located a few centimeters from the transmitting laptop, and the bitrate varies from 0.067 bit/sec to 2 bit/sec. This type of attack is less relevant on a standard workstation where there is a gap between the location of the internal hard drive and the chassis partition. Our method differs from [39] in the following respects:

Signal generation. We propose to generate the magnetic fields via the CPU which is available on virtually any computerized device today, including devices without magnetic HDDs. This method enables us to control the frequency of the transmissions independently for each CPU core, and hence to be able to use more complex modulation schemes.

Distance and bitrate. Our covert channel enables higher bitrates and transmission from greater distances, making it more feasible for the proposed attack scenarios.

Air-gaps and Faraday shielding. Our discussion focuses 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 protection and evasion in the context of covert channels and cyber-security measures. Table I summarizes the existing air-gap covert channels. Table II summarizes the differences between our work and existing work in the magnetic field.

The adversarial attack model requires running a malicious code in the targeted computer. In addition, there must also be a magnetic receiver hidden near the targeted system. The attack itself consists of four steps: (1) system infection, (2) receiver Implantation, (3) data gathering, and (4) exfiltration. **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 in recent years. Note that several APTs discovered in the last decade are capable of infecting air-gapped networks [40], e.g., Turla [41], RedOctober [42], and Fanny [43]. As a part of the targeted attack, the adversary may infiltrate

TABLE I
DIFFERENT TYPES OF AIR-GAP COVERT CHANNELS.

Type	Method
Electro-magnetic	AirHopper [4], [11] GSMem USBee [24] Funthenna [25]
Magnetic	ODINI (this paper) Myhayun (hard disk drive [39])
Acoustic	DiskFiltration (Hard disk noise) [28] Myhayun Fansmitter (computer fan noise) [27]
Thermal	BitWhisper [35]
Optical	Hard drive LED (LED-it-GO) [30] VisiSploit (invisible pixels) [32] Keyboard LEDs [30] Router LEDs [31]
Infrared (IR)	aIR-Jumper (security cameras & infrared) [34] Implanted infrared LEDs [33]

the air-gapped networks using social engineering, supply chain attacks, or insiders. **Receiver Implantation.** The receiver can be a hardware with a magnetic sensor hidden or implanted in close proximity to the target system. An example of such a hardware implant was presented in Snowden's leaked documents. The component, known as COTTONMOTH [44], is a USB connector implanted RF transceiver that attackers used to connect with the air-gapped networks. **Data gathering**. Having a foothold in the system, the malware starts retrieving interesting data for the attacker. The data might be textual data, encryption keys, credential tokens, or passwords. The data gathered is exfiltrated from a computer usually a PC workstation or server which contains the sensitive data to leak. **Exfiltration**. At the last phase of the attack, the malware starts the data leak by encoding the data and transmitting it via magnetic emanation generated from the computer. The transmissions may take place at predefined times or in response to some trigger infiltrated by the attacker. The leaked data is received by the nearby magnetic receiver and delivered to the attacker encrypted via standard networking (e.g., Wi-Fi). Note that although the described attack model is complicated, it is not beyond the capability of motivated and capable attackers. Advanced persistent threats coupled with sophisticated attack vectors such as supply chain attacks and human engineering have been shown to be feasible in the last decade. As a reward for these efforts, the attacker can get his/her hands on very valuable and secured information, which is out of reach of other types of covert channels.

D. Level of Isolation

The proposed covert channel is discussed in the context of highly isolated systems which are kept secluded

TABLE II
DIFFERENCES BETWEEN OUR WORK AND EXISTING WORK IN THE MAGNETIC FIELD

Work	Signal generation	Transmitters	Receiver	Max bitrate	Max distance
ODINI (this work)	CPU operations	PCs, servers, NUK, IoT and, embedded devices	Magnetic sensors	40 bit/sec	100 to 150 cm
Hard disk drive ([39])	Magnetic hard disk drive, I/O operations	Laptops with magnetic hard drive	Smartphones	2 bit/sec	4 to 12 cm
Near/far field magnetic, communication [38] [37]	Magnetic coils	Magnetic transmitters	Magnetic receivers	Tens to hundreds of bits, per second	Two meters to hundreds of,meters

with air-gaps and Faraday cages. However, the attack is also relevant in less restrictive environments as well, for example, computers that are not air-gapped or are not being kept in Faraday cages. In many cases, these computers are highly monitored in the network to detect a potential leakage of data and prevent it. In these cases, the attacker may choose to resort to a type of out-of-band air-gap covert channel, which is not monitored by existing security systems. Figure 1 illustrates a typical scenario for magnetic exfiltration, in which the magnetic receiver is located near highly secured air-gapped system, which is kept within Faraday shielding.
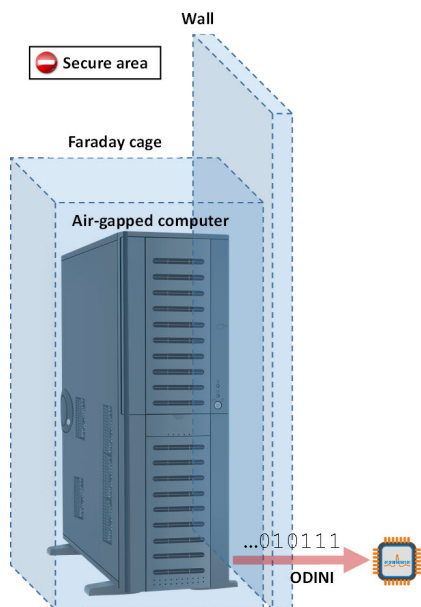

Fig. 1. An illustration of the magnetic covert channel (ODINI). Sensitive data is exfiltrated from the secured system, through air-gap and Faraday shielding.

III. 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 a 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 (1T) 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 ($1\mu T = 10^{-9}T$). Ampre'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 [45]. The magnetic flux density equation shows that the magnetic fields rapid decay is proportional to the inverse of the third power of the distance from the source:

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

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 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 magnetism [46]. 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 [46]. In practice, magnetic fields are used for the establishment of short

range wireless communication between close devices, a technique commonly referred to as near-field magnetic induction communication [38].

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 equal potential and prevents potential changes inside. Faraday shields may be small in size when protecting computer systems, or very large for protecting entire rooms and laboratories [19]. 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 [13]. 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 [47].

C. Magnetic Fields and Metal Shielding

The propagation of electromagnetic and magnetic radiation in conducting mediums such as concrete is better in the low frequencies [45], [48]. 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 signal in 100KHz would require an antenna that is more than 3km 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 [37], [45]. In the proposed covert channel, we generate magnetic fields at frequencies lower than 50Hz. It is known that low frequency magnetic fields have a low impedance and are difficult to block with metal shielding, since this would require very thick metal surfaces [49]. The following analysis shows that magnetic fields at low frequencies can bypass the metal shields of computer
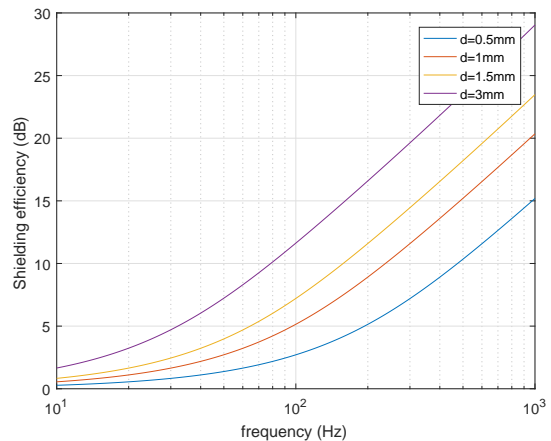

Fig. 2. Shielding efficiency of the closed metal cube shell at a thickness of 0.5mm to 3mm.

chassis and Faraday cages. Figure 2 shows the attenuation of a magnetic field (the reduction in magnitude of magnetic field strength) given a thin walled cube metal shield, based on the shielding approximation formulas [50]. The field attenuation is measured in decibels (dB) and is equal to $20log\frac{E1}{E2}$, 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 low frequencies (< 50Hz), as the magnetic attenuation is as 5dB at most. The results show that low frequency magnetic fields penetrate the typical computer chassis and Faraday shielding.

IV. TRANSMISSION

In this section we describe the signal generation algorithm and present the data modulation schemas 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. The CPU is one of the greatest consumers of power in the motherboard. Since modern CPUs are energy efficient, the momentary workload of the CPU directly affects the dynamic changes in its power consumption [51]. 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. We developed a fine-grained approach, in which we control the workload of each of the CPU core 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. Choosing which cores to operate on at a given time, allows us to use only the currently available cores, that is, cores which are not utilized by other processes. This way, the transmission activity wont interrupt other active processes in the system. This is important for the usability of the computer and the stealth of the covert channel.

2) Controlling the signal strength. By using different numbers of cores for the transmission, we can control the strength of the magnetic field (e.g., fewer cores consume less power), and hence the amplitude of the carrier wave. This allows us to employ amplitude based 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).

To generate a carrier wave at frequency f_c in one or more cores, we control the utilization of the CPU at a frequency correlated to f_c. To that end, n worker threads are created where each thread is bound to a specific core. To generate the carrier wave, each worker thread overloads its core at a frequency f_c repeatedly alternating between applying a continuous workload on its core for a time period of $1/2f_c$ (full power consumption) and putting its core in an idle state for a time period of $1/2f_c$ (no power consumption).

This operation is illustrated in Figure 3, which depicts a system with two worker threads. Threads $T3$ and $T4$ are bound to cores $C3$, $C4$, respectively. Note that cores $C1$ and $C2$ dont participate in this transmission. When the worker threads $T3$ and $T4$ start, they receive the required carrier frequency f_c and the stream of bits
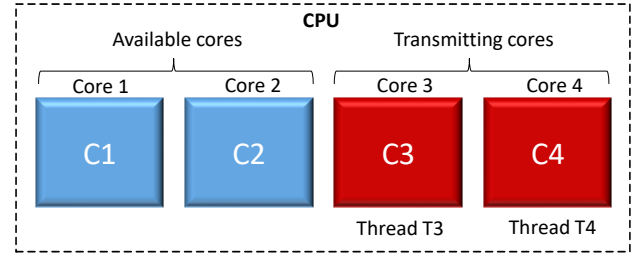

Fig. 3. Signal generation using two threads and two CPU cores.

Algorithm 1 WorkerThread (iCore, freq, nCycles0, nCycles1)

1: $bindThreadToCore(iCore)$
2: $cycle\_ms \leftarrow 1000/freq$
3: $half\_cycle\_ms \leftarrow cycle\_ms * 0.5$
4: **while** $!endTransmission()$ **do**
5:     **if** thread.data[i] == 0 **then**
6:         $sleep(nCycles0 * cycle\_ms)$
7:     **else**
8:         **for** $j \leftarrow 0\ to\ nCycles1$ **do**
9:             $busywait(half\_cycle\_ms)$
10:             $sleep(half\_cycle\_ms)$
11:         **end for**
12:     **end if**
13:     $i + +$
14: **end while**

---

to transmit. The basic operation of a worker thread is described in algorithm 1.

A worker thread receives the core to be bound to (iCore) and the carrier frequency (freq). It also receives the number of cycles for the modulation of logical '0' (nCycles0) and the number of cycles for the modulation of logical '1' (nCycles1). Note that the cycle time is derived from the frequency of the carrier wave (lines 2-3). The threads main function iterates on the array of bits to transmit. In the case of logical '0' it sleeps for nCycles0 cycles (line 6). In the case of logical '1' it repeatedly starts and stops the workload of the core at the carrier frequency freq for nCycles1 cycles (lines 8-11). 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. Based on the algorithm, we implemented

---

Algorithm 2 busywait(ms)

1: $T1 \leftarrow getCurrentTime()$
2: **while** $(getCurrentTime() - T1 < ms)$ **do ;**
3: **end while**

---

a transmitter for Linux Ubuntu (version 16.04, 64 bit). We used the sched_setaffinity system call to bind each thread to a CPU core. The affinity is the thread level attribute that is configured independently for each worker thread. To synchronize the initiation and termination of the worker threads, we used the thread mutex objects with pthread_mutex_lock() and pthread_mutex_unlock() [52]. For thread sleeping we used the sleep() system call [53]. Note that the precision of sleep() in milliseconds is sufficient given the low frequencies of the carrier waves (e.g., a sleeping cycle of 600ms is required for a 50Hz carrier wave).

1) Stealth: The transmitting program leaves only a small footprint in the memory, making its presence easier to hide from AVs. At the OS level, the transmitting program requires no special or elevated privileges (e.g., root or admin), and hence can be initiated from an ordinary user space process. The transmitting code mainly consists of basic CPU operations such as busy loops, which do not expose malicious behaviors, making it highly evasive from automated analysis tools.

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 a limited amount of control of the carrier waves 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: On-off keying (OOK), amplitude-shift keying (ASK), and frequency-shift keying (FSK). We also implement the more efficient orthogonal frequency-division multiplexing (OFDM) modulation scheme. Note that the limited control of the amplitude of the generated carrier wave, also affects the control of the signals phase. Thus, we did not use phase-based modulations (e.g., PSK) for this covert channel. In the following sections, we describe each of the four modulations used. We denote the number of cores available for the transmission as N_c.

1) On-Off Keying : In OOK modulation, 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' (Table III). 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 and bit framing which are described later.

TABLE III
ON-OFF KEYING

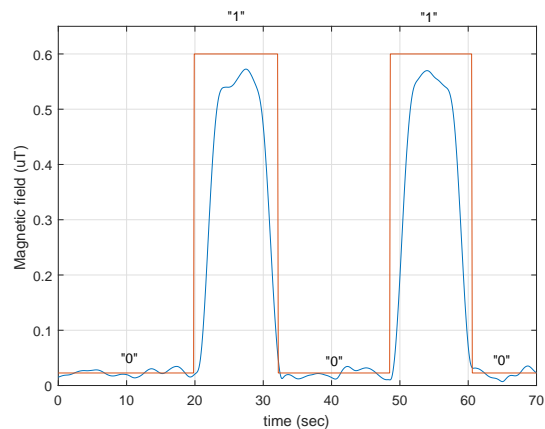| Symbol	Carrier wave	# of cores
0	Present	$n \leq N_c$
1	Absent	$n \leq N_c$

Fig. 4. The waveform of a binary sequence ('01010') modulated with OOK.

Figure 4 shows the waveform of a binary sequence ('01010') modulated with OOK and transmitted from a desktop PC with four cores. A magnetic sensor is located 20 cm away from the transmitting computer. The noise level in this case is 0.02mT, and the carrier wave amplitude is 0.58mT (SNR of 29dB).

2) 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 signal strength (amplitude) 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 IV 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 5 shows the waveform of a binary sequence ('11100100') modulated with four level ASK and transmitted from a desktop PC with four cores. A magnetic sensor is located at a distance of 20 cm from the transmitting computer. In this case, the symbols '00', '01', '10,' and '11' are represented by four level of amplitudes A_0, A_1, A_2, and A_3, and generated by one, two, three, and four cores, respectively.

TABLE IV
AMPLITUDE-SHIFT KEYING

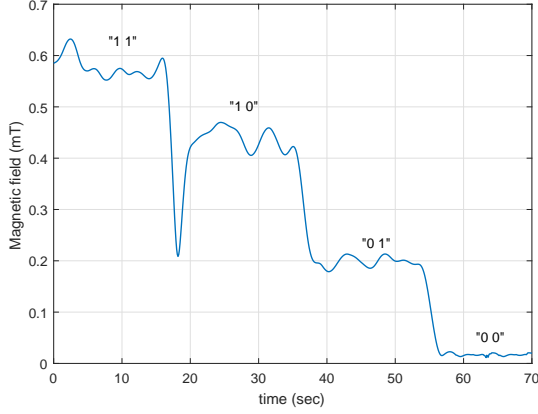| Symbol	Amplitude	# of cores
00	A_0	0 (or 1)
01	A_1	2
10	A_2	3
11	A_3	4

Fig. 5. ASK modulation of 111000100 with four amplitudes.

3) Frequency-Shift Keying: In frequency-shift keying (FSK) the data is represented by a change in the frequency of a carrier wave. Recall that the transmitting code can determine the frequency of the signal by setting the cycle time in the signal generation algorithm. In FSK, each frequency represents a different symbol. Table V presents a case in which we encode the two symbols '0' and '1' with two frequencies F_0, and F_1 .

Figure 6 shows the time-frequency spectrogram of a binary sequence ('010101010') modulated with three frequencies FSK as transmitted from a PC with four cores. In this modulation, the frequencies 3Hz, 7Hz and 13Hz have been used to encode 0, 1 and 01 respectively. A magnetic sensor is located at a distance of 20 cm from the transmitting computer.

4) 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 in a range of 0-50Hz. In each sub-carrier, we used OOK to modulate the data. Note that since the sub-

TABLE V
FREQUENCY-SHIFT KEYING

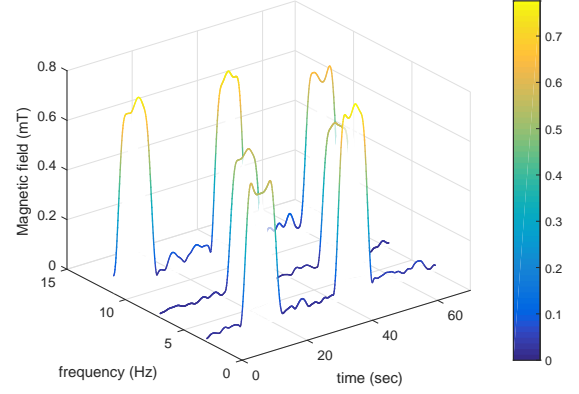| Symbol	Amplitude	# of cores
0	F_0	$n \leq N_c$
1	F_1	$n \leq N_c$

Fig. 6. FSK modulation of a binary sequence ('010101010') with three frequencies (3Hz, 7Hz and 13Hz)

carriers signals are generated in parallel, the maximal number of sub-carriers is equal to the number of cores available for the transmissions nN_c. For example, in the case of two available cores, we define two sub-carriers (e.g., 7Hz and 13Hz) for the transmission of the data stream.

Figure 7 presents the binary sequence ('1101111011') modulated with OFDM with two sub-carriers as transmitted from a PC with four cores. In this modulation, 7Hz and 13Hz have been used to encode the symbols '00', '01', '10' and '11'. A magnetic sensor is located 20 cm from the transmitting computer.

C. Bit-Framing

We transmit the data in small packets composed of a preamble, a payload, and a parity bit.

- Preamble. Like most air-gap covert channels, the unidirectional communication means that the receiver cannot establish a handshake with the transmitter, and hence cannot determine or set the channel parameters before the transmission. To solve this, 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 properties of the channel, such as the carrier wave frequency and amplitude. In addition, the preamble allows the receiver to detect the beginning of a transmission and synchronize itself.
- 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
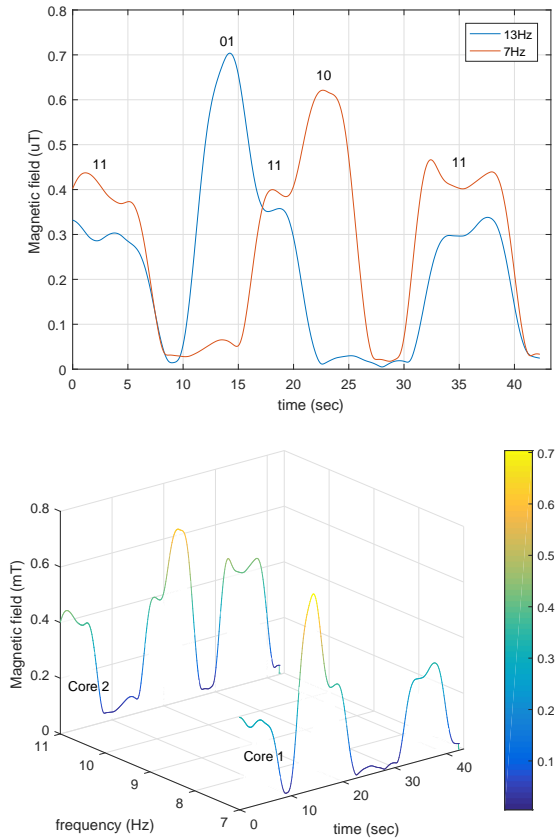
Fig. 7. The waveform and spectrogram of the binary sequence ('1101111011') modulated with OFDM and two sub-carriers (7Hz and 11Hz).

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.

V. ANALYSIS & EVALUATION

In this section, we present an analysis and evaluation of the proposed covert channel. Note that although near-field and far-field magnetic communication are known research topics[36], [38], such magnetic communication requires specialized and dedicated magnetic transceivers, which are not available in the case of the proposed covert channel. 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 using Linux Ubuntu version 16.04 for 64-bit. A detailed list of the computers is provided in Table VI.

Hyper-Threading. Note that modern Intel CPUs support the Hyper-Threading technology [54] . 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 system's 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.

2) Recevier : For the reception, we used the Honeywell HMR2300 magnetic sensor [55]. This is a digital magnetometer which is capable of sampling the strength and direction of a 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. The output is converted to 16-bit digital values using an internal analog-to-digital converter. The sensor resolution of approximately 70nT (nano-teslas) and the sampling rate is up to 154 samples per second.

3) Measurement Setup: The measurement setup is shown in Figure fig. 8. The Honeywell HMR2300 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 the system driven by a LabVIEW data flow visual programming language.
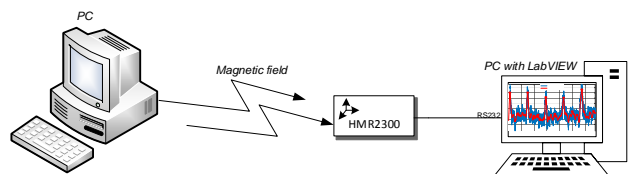

Fig. 8. The measurement setup

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 9

TABLE VI
THE COMPUTERS USED IN THE EXPERIMENTAL SETUP

#	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 threads)
2	PC-2,(Desktop PC)	Lenovo desktop PC	Panda L-IQ45	Intel Core Quad-Q9550 CPU @ 2.83GHz	4 (4 threads)
3	Laptop	HP ProBook 650 G2	Intel	Intel Core i5-Q6200U CPU @ 2.4GHz	2 (4 threads)
4	Server	IBM System x3500 M4	Intel C602J	Intel Xeon CPU E5-2620	12 (24 threads)
5	NUK,(small form factor)	Intel NUK. Lenovo ThinkCentre M93p	Intel Q87 express chipset for ThinkCentre M93/M93p	Intel Core i7 -4785T	4 (8 threads)
6	IOT,(IoT/,embedded)	Raspberry Pi 3	Raspberry Pi 3 model B V1.2	Quad Core Broadcom BCM2837 64-bit ARMv8, processor Cortex A53	4

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 20 cm from the transmitting computers showed a gradual increase when an increasing number 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.
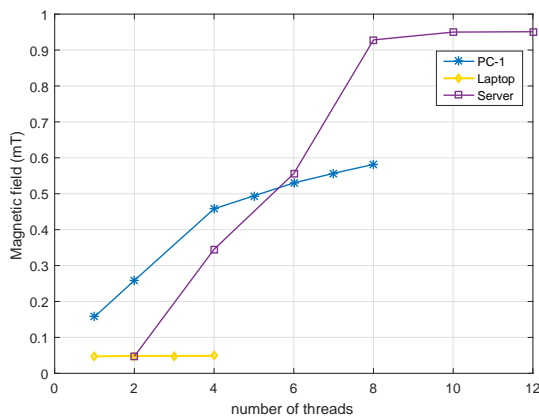
sured at various distances from five transmitters. Note that Figure 9 shows the magnetic field in a logarithmic scale. Using the HMR3200 sensor, the magnetic signals were received at a maximal distance of 50 cm for the laptop, 100 cm for the desktop PCs and small form factor PC, and 150 cm for the server. Given the resolution of the HMR2300 sensor and the signal to noise ratio (SNR) levels, we stopped the measurements at a field power of 10^{-2} mT. Note that the reception of the magnetic signals at greater distances requires more sensitive sensors such as search coil or fluxgate magnetometers [56].
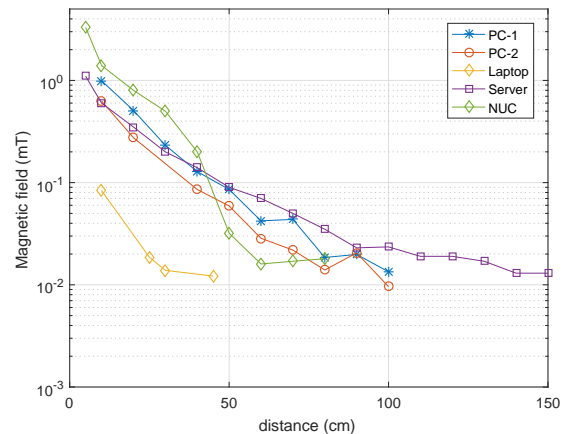

Fig. 10. A 20Hz signal generated by five computers (PC-1, PC-2, laptop, server, and NUK) measured at distances of 0 to 150 cm.

Figure 11 depicts the reception of a 4.7Hz signal from PC-1 at a distance of 100 cm. It shows the measured signal in the frequency domain, between 0 and 50Hz. 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 is shows the background magnetic noise,

Fig. 9. Measurement of the magnetic signal generated by different number of threads on three computers: PC-1, laptop, and server

2) Distance : The strength of a magnetic field decreases fast, inversely proportional to the third power of the distance from the magnetic source. Figure 10 and Table VII show the strength of magnetic signals as mea-

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

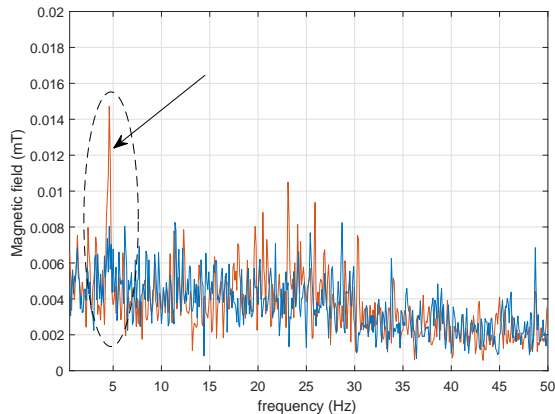|	10 cm	20 cm	40 cm	60 cm	80 cm	100cm	120cm	140cm	150cm
PC-1	0.99 mT	0.51 mT	0.13 mT	0.042 mT	0.019 mT	0.013 mT	-	-	-
PC-2	0.63 mT	0.27 mT	0.059	0.022 mT	0.01 mT	0.009 mT	-	-	-
Laptop	0.084 mT	0.019 mT	0.012 mT	-	-	-	-	-	-
Server	0.6 mT	0.35 mT	mT 0.14	0.07 mT	0.035 mT	0.026 mT	0.019 mT	0.017 mT	0.013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

Fig. 11.   A 4.7Hz signal transmitted from Faraday shielded PC-1 as received from a distance of one 100 cm away.

Fig. 12.   The channel capacity of five transmitters based on the SNR measured at a range of distances.

while the red line shows the sampled signal. A signal strength of 0.014mT is observed at 4.7Hz, which is above the average noise, which is less than 0.008 mT in this band (SNR of 4.8dB).

C. Channel Capacity

We calculated the maximal bitrate of the covert channel based on the Shannon-Hartley channel capacity limit. Figure 10 shows the channel capacity given the HMR2300 receiver for PC-1, PC-2, laptop, server, and NUK transmitters. In our case, the bandwidth (B) is 50Hz given the sampling rate of the HMR2300 sensor and the low frequencies we used for the transmissions. The S and N were calculated based on the SNR measurements taken for each of the computers at distances of 0 to 150 cm. As shown in Figure 11, for desktop and server computers, the channel capacity varies from 300 bit/sec to 30 bit/sec depending on the distance from the computer. The laptop as a transmitter yields a lower channel capacity due to the lower magnetic signals it generates.

D. Data Transfer

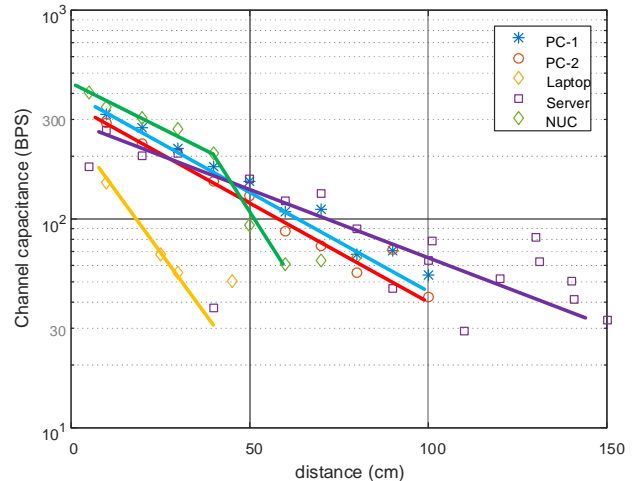The channel capacity represents the upper theoretical limits of a communication channel. The actual bitrate is usually lower than the channel capacity and is determined by the modulation scheme and the quality of the transmitter and receiver used. We measured the bit error rate (BER) of PC-1, the server, 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.

The results, presented in Table VIII, 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 13 shows the waveforms of a binary sequence ('10101110110101010111') encoded in OOK, as transmitted from the server computer. The data was transmitted at a speed of 5 bit/sec and received at distances of

TABLE VIII
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%	20%
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%	-	-	-	-	-
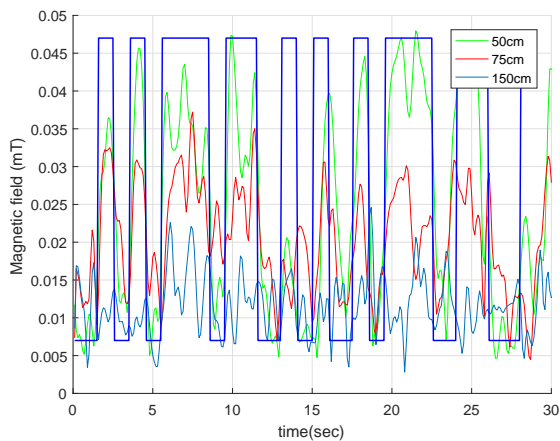
Fig. 13. . Server data measurements for distances of 50, 75, and 150 cm (BER of 0 .

Fig. 14. The waveform of 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 desktop/server virtualization systems and 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 virtualization layer caused interruptions or delays which may affect the signal generation. Figure 15 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 a VMWare virtual machine. Both signals represent the transmission of the alternating sequence (101010) using OOK modulation. 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 SNR of 15dB. More specifically, we experienced no time delay or reduction in the power

50 cm, 75 cm, and 150 cm away, with an SNR of 10dB, 8dB and 4.4dB respectively. It is possible to increase the effective bitrate further by employing advanced signal processing algorithms or by using magnetometers with increased sensitivity and resolution, however we leave the exploration of both of these directions to future work in this field.

1) Embedded/IoT Device: Embedded devices usually consume just a small amount of power, hence emitting lower magnetic fields. Our experiments show that the proposed covert channel also works for such low power devices, when the magnetic sensor is in close proximity of the device. Figure 14 shows the waveform of 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% with a SNR of 15dB.
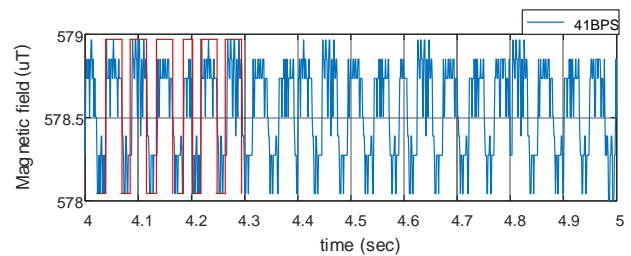
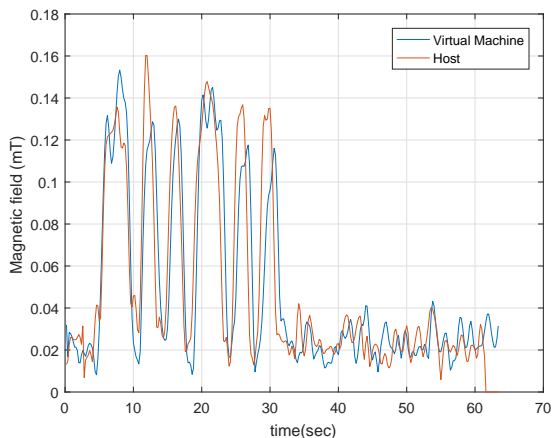of the signal when it was generated from the VM.

Fig. 15. The waveforms of two signals transmitted from PC-1 (VM/Host).

We also investigated the feasibility of controlling the signal strength from virtual machines by using different numbers of cores. Figure 16 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 the host-based transmissions. In the context of the communication channel, it allows the attacker to use amplitude modulations from VMs.
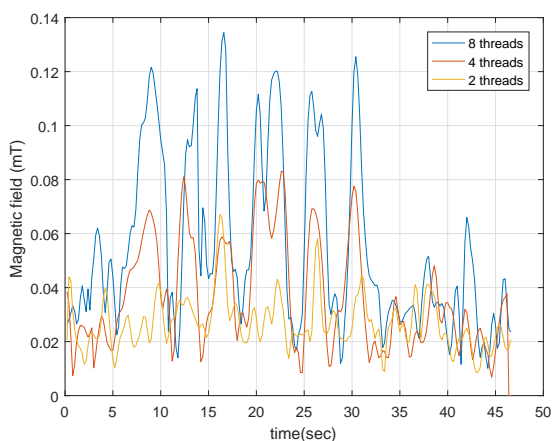

Fig. 16. The waveform of a signal generated from a virtual machine using 2, 4, and 8 threads.

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 activity 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 in desktop PCs:

1) Idle. The system was idle and only the default processes were running in the background.
2) Word processing. The LibreOffice Writer [57] was open, and the user typed a document.
3) Video playing. The VLC media player [58] was playing an HD video clip.
4) Backup. The Linux rsync [59] command was performing a backup of local folders in the HDD.
5) CPU intensive calculations. The Linux mathoprimes [60] was performing the calculations of big prime numbers.

Table IX summarizes the SNR measured at a distance of 20 cm from the transmitting computer for each of the five workloads. We used eight threads for the transmission of an alternating bit sequence ('10101010') using OOK modulation. 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 35 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.

VI. 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, intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) that continuously trace the activities of a computers processes and try to detect malicious operations; in the case of a magnetic covert channel, a thread (or group of threads) that abnormally regulates the CPU workload would be

TABLE IX
SNR WITH VARIOUS WORKLOADS

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

tagged as suspicious. However, many types of applications use working threads that affect the processors workload, and therefore, such a detection approach would likely suffer from a high rate of false alarms. Another problem in the runtime detection approach is that the signal generation code only involves simple, non-privileged CPU operations (e.g., busy loops), without requiring special instructions or specialized API calls. Tracing non-privileged CPU operations at runtime necessitates entering the processes to a step-by-step mode, which can severely degrade system performance [15]. Software based detection also suffers from an inherent weakness in that they can be easily bypassed by malware using a wide range of evasion techniques. Another approach is to detect the covert channel externally, by monitoring the magnetic field in the area of the computer. The magnetic field measured is continuously processed to find hidden transmissions or deviations from the standards. Note that finding anomalies in magnetic and electromagnetic spectrums may also suffer from a high rate of false positives [61], [62].

B. Prevention

There are three different approaches that can be used to prevent attackers from establishing a magnetic covert channel: shielding, jamming, and zoning. 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 [63]. 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 [49]. 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 [64]. Signal Jamming. Signal jamming is commonly used to mitigate electromagnetic and acoustic covert channels [65]. 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) [66]. 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 [67]. 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). 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 NATO standards SDIP-27 and SDIP-28 define separated zones in which electronic equipment are allowed [68]. 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 X.

VII. CONCLUSION

This paper presents a new type of covert channel based on magnetic fields generated by the computer CPU. This

TABLE X
LIST OF COUNTERMEASURES

Countermeasure	Description	Cons
Malicious activity detection (software)	Detect the transmitting threads	False positives, Can be bypassed (e.g., by rootkits)
Magnetic activity detection (hardware)	Detect abnormal magnetic field activities	False positives, Expensive
Ferromagnetic shielding (mu-metal)	Shielding with Ferromagnetic material	Expensive
Magnetic field jammer (hardware)	Jamming the transmissions with magnetic field generator	Expensive
Field cancellation	Producing the counter magnetic fields	Expensive
Random workload generator	Starting random transmissions which jam the transmitter,signals	Degrades system performance, Can be disabled (e.g., by rootkits)
Zone separation	Banning magnetic receivers/electronic equipment	Expensive (in terms of physical,space)

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 easily penetrate through metals. This makes our covert channel possible even in a constrained environment where the computers are enclosed within Faraday shielding. We present scientific background and explain the characteristics of magnetic fields and the signal generation process. 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 and devices. We also propose different types of defensive countermeasures to detect and prevent this threat. Our results show that data can be exfiltrated from air-gap computers via low frequency magnetic fields at bitrates of 1-40 bit/sec. Notably, this type of covert channel can evade Faraday shielding protection, where conventional electromagnetic covert channels fail.

REFERENCES

[1] MacAskill E, Thielman S, Oltermann P. WikiLeaks publishes' biggest ever leak of secret CIA documents'. The Guardian Retrieved. 2017;26.

[2] Zetter K. Sony got hacked hard: What we know and dont know so far. Wired, January. 2014;.

[3] Thielman S. Yahoo hack: 1bn accounts compromised by biggest data breach in history. The Guardian, Dezember. 2016;.

[4] Guri M, Monitz M, Elovici Y. Bridging the Air Gap between Isolated Networks and Mobile Phones in a Practical Cyber-Attack. ACM Transactions on Intelligent Systems and Technology (TIST). 2017;8(4):50.

[5] Byres E. The air gap: SCADA's enduring security myth. Communications of the ACM. 2013;56(8):29–31.

[6] Classified United States website - Wikipedia;. (Accessed on 12/03/2017). https://en.wikipedia.org/wiki/Classified_United_States_website.

[7] Maybury M, Chase P, Cheikes B, Brackney D, Matzner S, Hetherington T, et al. Analysis and detection of malicious insiders. MITRE CORP BEDFORD MA; 2005.

[8] Trump, Putin, and the New Cold War - The New Yorker;. (Accessed on 12/03/2017). https://www.newyorker.com/magazine/2017/03/06/trump-putin-and-the-new-cold-war.

[9] Abraham S, Chengalur-Smith I. An overview of social engineering malware: Trends, tactics, and implications. Technology in Society. 2010;32(3):183–196.

[10] Wikileaks: CIA uses 'Brutal Kangaroo' toolkit to hack air-gapped networks;. (Accessed on 12/03/2017). https://www.theinquirer.net/inquirer/news/3012499/-wikileaks-cia-uses-brutal-kangaroo-toolkit-to-hack-air-gapped-networks.

[11] Guri M, Kedma G, Kachlon A, Elovici Y. AirHopper: Bridging the air-gap between isolated networks and mobile phones using radio frequencies. In: Malicious and Unwanted Software: The Americas (MALWARE), 2014 9th International Conference on. IEEE; 2014. p. 58–67.

[12] Kuhn MG, Anderson RJ. Soft Tempest: Hidden Data Transmission Using Electromagnetic Emanations. In: Information hiding. vol. 1525. Springer; 1998. p. 124–142.

[13] Kuhn MG. Compromising emanations: eavesdropping risks of computer displays. University of Cambridge; 2002.

[14] Vuagnoux M, Pasini S. Compromising Electromagnetic Emanations of Wired and Wireless Keyboards. In: USENIX security symposium; 2009. p. 1–16.

[15] Guri M, Kachlon A, Hasson O, Kedma G, Mirsky Y, Elovici Y. GSMem: Data Exfiltration from Air-Gapped Computers over GSM Frequencies. In: USENIX Security Symposium; 2015. p. 849–864.

[16] Hoeft LO, Hofstra JS. Measured electromagnetic shielding performance of commonly used cables and connectors. IEEE Transactions on Electromagnetic Compatibility. 1988;30(3):260–275.

[17] EMP - EMI shielded racks;. (Accessed on 12/03/2017). https://hollandshielding.com/EMP-EMI-shielded-racks.

[18] EMF Shielding Devices for Computers and TVs;. (Accessed on 12/03/2017). https://www.lessemf.com/computer.html.

[19] www.comtest.eu/products/rf-shielded-rooms-doors/rf-shielded-rooms.html;. (Accessed on 12/03/2017). http://www.comtest.eu/products/rf-shielded-rooms-doors/rf-shielded-rooms.html.

[20] Giani A, Berk VH, Cybenko GV. Data exfiltration and covert channels. In: Defense and Security Symposium. International Society for Optics and Photonics; 2006. p. 620103–620103.

[21] Murdoch SJ, Lewis S. Embedding covert channels into TCP/IP. In: Information hiding. vol. 3727. Springer; 2005. p. 247–261.

[22] Zander S, Armitage G, Branch P. A survey of covert channels and countermeasures in computer network protocols. IEEE Communications Surveys & Tutorials. 2007;9(3):44–57.

[23] Tempest for Eliza;. (Accessed on 12/03/2017). http://www.erikyyy.de/tempest/.

[24] Guri M, Monitz M, Elovici Y. USBee: Air-gap covert-channel via electromagnetic emission from USB. In: Privacy, Security and Trust (PST), 2016 14th Annual Conference on. IEEE; 2016. p. 264–268.

[25] funtenna GitHub;. (Accessed on 12/03/2017). https://github.com/funtenna.

[26] Hanspach M, Goetz M. On covert acoustical mesh networks in air. arXiv preprint arXiv:14061213. 2014;.

[27] Guri M, Solewicz Y, Daidakulov A, Elovici Y. Fansmitter: Acoustic Data Exfiltration from (Speakerless) Air-Gapped Computers. arXiv preprint arXiv:160605915. 2016;.

[28] Guri M, Solewicz Y, Daidakulov A, Elovici Y. Acoustic Data Exfiltration from Speakerless Air-Gapped Computers via Covert Hard-Drive Noise (DiskFiltration). In: European Symposium on Research in Computer Security. Springer; 2017. p. 98–115.

[29] Loughry J, Umphress DA. Information leakage from optical emanations. ACM Transactions on Information and System Security (TISSEC). 2002;5(3):262–289.

[30] Guri M, Zadov B, Elovici Y. In: Polychronakis M, Meier M, editors. LED-it-GO: Leaking (A Lot of) Data from Air-Gapped Computers via the (Small) Hard Drive LED. Cham: Springer International Publishing; 2017. p. 161–184. Available from: https://doi.org/10.1007/978-3-319-60876-1_8.

[31] Guri M, Zadov B, Daidakulov A, Elovici Y. xLED: Covert Data Exfiltration from Air-Gapped Networks via Router LEDs. arXiv preprint arXiv:170601140. 2017;.

[32] Guri M, Hasson O, Kedma G, Elovici Y. An optical covert-channel to leak data through an air-gap. In: Privacy, Security and Trust (PST), 2016 14th Annual Conference on. IEEE; 2016. p. 642–649.

[33] Lopes AC, Aranha DF. Platform-agnostic Low-intrusion Optical Data Exfiltration. In: ICISSP; 2017. p. 474–480.

[34] Guri M, Bykhovsky D, Elovici Y. aIR-Jumper: Covert Air-Gap Exfiltration/Infiltration via Security Cameras & Infrared (IR). arXiv preprint arXiv:170905742. 2017;.

[35] Guri M, Monitz M, Mirski Y, Elovici Y. Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations. In: Computer Security Foundations Symposium (CSF), 2015 IEEE 28th. IEEE; 2015. p. 276–289.

[36] Sojdehei JJ, Wrathall PN, Dinn DF. Magneto-inductive (MI) communications. In: OCEANS, 2001. MTS/IEEE Conference and Exhibition. vol. 1. IEEE; 2001. p. 513–519.

[37] Through-The-Earth Two-Way Emergency Wireless Communications for Mine Industry Safety;. (Accessed on 12/03/2017). http://www.teslasociety.ch/info/magnetlink/2.pdf.

[38] Bansal R. Near-field magnetic communication. IEEE Antennas and Propagation Magazine. 2004;46(2):114–115.

[39] Matyunin N, Szefer J, Biedermann S, Katzenbeisser S. Covert channels using mobile device's magnetic field sensors. In: Design Automation Conference (ASP-DAC), 2016 21st Asia and South Pacific. IEEE; 2016. p. 525–532.

[40] Industrial Defence In-Depth, Kaspersky Lab;. (Accessed on 12/03/2017). https://www.sans.org/summit-archives/file/summit-archive-1493412875.pdf.

[41] The Epic Turla (snake/Uroburos) attacks — Virus Definition — Kaspersky Lab;. (Accessed on 12/03/2017). https://www.kaspersky.com/resource-center/threats/epic-turla-snake-malware-attacks.

[42] ZAO K. Red october diplomatic cyber attacks investigation;.

[43] A Fanny Equation: "I am your father, Stuxnet" - Securelist;. (Accessed on 12/03/2017). https://securelist.com/a-fanny-equation-i-am-your-father-stuxnet/68787/.

[44] COTTONMOUTH-III: NSA Exploit of the Day - Schneier on Security;. (Accessed on 12/03/2017). https://www.schneier.com/blog/archives/2014/03/cottonmouth-iii.html.

[45] Kodali VP. Engineering electromagnetic compatibility: principles measurements technologies and computer models. Institute of Electrical and Electronics Engineers; 2001. .

[46] Sadiku MN. Elements of electromagnetics. Oxford university press; 2014.

[47] Bindar V, Popescu M, Craciunescu R. Aspects of electromagnetic compatibility as a support for communication security based on TEMPEST evaluation. In: Communications (COMM), 2014 10th International Conference on. IEEE; 2014. p. 1–4.

[48] Holloway CL, Hill DA, Dalke RA, Hufford GA. Radio wave propagation characteristics in lossy circular waveguides such as tunnels, mine shafts, and boreholes. IEEE transactions on antennas and propagation. 2000;48(9):1354–1366.

[49] EMC for Systems and Installations Part 4 - Filtering and Shielding;. (Accessed on 12/03/2017). http://www.compliance-club.com/archive/keitharmstrong/systems_installations4.html.

[50] Kelha V, Pukki J, Peltonen R, Penttinen A, Ilmoniemi R, Heino J. Design, construction, and performance of a large-volume magnetic shield. IEEE Transactions on Magnetics. 1982;18(1):260–270.

[51] von Kistowski J, Block H, Beckett J, Spradling C, Lange KD, Kounev S. Variations in cpu power consumption. In: Proceedings of the 7th ACM/SPEC on International Conference on Performance Engineering. ACM; 2016. p. 147–158.

[52] pthread_mutex_lock(3): lock/unlock mutex - Linux man page;. (Accessed on 12/03/2017). https://linux.die.net/man/3/pthread_mutex_lock.

[53] sleep(3) - Linux manual page;. (Accessed on 12/03/2017). http://man7.org/linux/man-pages/man3/sleep.3.html.

[54] Marr D, Binns F, Hill D, Hinton G, Koufaty D, et al. Hyper-threading technology in the netburst® microarchitecture. 14th Hot Chips. 2002;.

[55] Honeywell HMR2300;. (Accessed on 12/03/2017). https://aerocontent.honeywell.com/aero/common/documents/myaerospacecatalog-documents/Missiles-Munitions/HMR2300.pdf.

[56] Auster H, Glassmeier K, Magnes W, Aydogar O, Baumjohann W, Constantinescu D, et al. The THEMIS fluxgate magnetometer. In: The THEMIS Mission. Springer; 2009. p. 235–264.

[57] Home — LibreOffice - Free Office Suite - Fun Project - Fantastic People;. (Accessed on 01/11/2018). https://www.libreoffice.org/.

[58] Official download of VLC media player, the best Open Source player - VideoLAN;. (Accessed on 01/11/2018). https://www.videolan.org/vlc/index.html.

[59] rsync(1) - Linux man page;. (Accessed on 01/14/2018). https://linux.die.net/man/1/rsync.

[60] Ubuntu Manpage: matho-primes - generate consecutive prime numbers;. (Accessed on 01/11/2018). http://manpages.ubuntu.com/manpages/zesty/man1/matho-primes.1.html.

[61] Carrara B. Air-Gap Covert Channels. Université d'Ottawa/University of Ottawa; 2016.

[62] Goher SZ, Javed B, Saqib NA. Covert channel detection: A survey based analysis. In: High capacity optical networks

and enabling technologies (HONET), 2012 9th international conference on. IEEE; 2012. p. 057–065.

[63] ter Brake HJ, Wieringa H, Rogalla H. Improvement of the performance of a mu-metal magnetically shielded room by means of active compensation (biomagnetic applications). Measurement Science and Technology. 1991;2(7):596.

[64] Yashchuk VV, Lee SK, Paperno E. In: Budker D, Jackson Kimball DFE, editors. Magnetic shielding. Cambridge University Press; 2013. p. 225248.

[65] Wilhelm M, Martinovic I, Schmitt JB, Lenders V. WiSec 2011 demo: RFReact—a real-time capable and channel-aware jamming platform. ACM SIGMOBILE Mobile Computing and Communications Review. 2011;15(3):41–42.

[66] Schloeder-EMV: Magnetic Field Generator and Analyzer;. (Accessed on 12/03/2017). http://www.schloeder-emv.de/en/emc-products/emc-test-and-measurement-system/emc-generators-measurement-systems/magnetic-field-generator-and-analyzer.html.

[67] Okazaki Y, Yanase S, Sugimoto N. Active magnetic shielding with magneto-impedance sensor. International Journal of Applied Electromagnetics and Mechanics. 2001;13(1-4):437–440.

[68] Anderson R. Emission security. Security Engineering,. 2008;p. 523–546.