Physiological Information Leakage: A New Frontier in Health Information Security

Arsalan Mohsen Nia, Susmita Sur-Kolay, *Senior Member, IEEE,* Anand Raghunathan, *Fellow, IEEE,* and Niraj K. Jha, *Fellow, IEEE*

This paper can be cited as: A. Mohsen Nia and N. K. Jha, "Physiological information leakage: A new frontier in health information security," in IEEE Trans. Emerging Topics in Computing, vol. 4, no. 3, pp. 321-334, 2016

The latest version of this manuscript is available on http://ieeexplore.ieee.org/document/7273876/

***Abstract*—Information security has become an important concern in healthcare systems, owing to the increasing prevalence of medical devices and the growing use of wearable and mobile computing platforms for health and lifestyle monitoring. Previous work in the area of health information security has largely focused on attacks on the wireless communication channel of medical devices, or on health data stored in online databases.**

In this work, we pursue an entirely different angle to health information security, motivated by the insight that the human body itself is a rich source (acoustic, visual, and electromagnetic) of data. We propose a new class of information security attacks that exploit *physiological information leakage*, i.e., various forms of information that naturally leak from the human body, to compromise privacy. As an example, we demonstrate attacks that exploit acoustic leakage from the heart and lungs.

Next, the medical devices deployed within or on our bodies also add to natural sources of physiological information leakage, thereby increasing opportunities for attackers. Unlike previous attacks on medical devices, which target the wireless communication to/from them, we propose privacy attacks that exploit information leaked by the very operation of these devices. As an example, we demonstrate how the acoustic leakage from an insulin pump can reveal important information about its operation, such as the duration and dosage of insulin injection. Moreover, we show how an adversary can estimate blood pressure (BP) by capturing and processing the electromagnetic radiation of an ambulatory BP monitoring device.

***Index Terms*—Healthcare, information leakage, information security, medical devices, privacy.**

I. Introduction

Implantable and wearable medical devices (IWMDs) promise to transform healthcare, by enabling diagnosis, monitoring, and therapy for a wide range of medical conditions and by facilitating improved and healthier lifestyles. Rapid advances in electronic devices are revolutionizing the capabilities of IWMDs [1]. New generations of IWMDs feature increased functional complexity, programmability, and wireless connectivity to body-area networks (BANs). Standardized communication protocols, such as Bluetooth [2] and ZigBee [3], are opening up new opportunities for providing low-power and reliable communication to IWMDs. These features facilitate convenient collection of medical data and personalized tuning of therapy through communication between different IWMDs and an external device (e.g., smartphone or clinical diagnostic equipment).

Advances in IWMDs have, unfortunately, also greatly increased the possibility of security attacks against them. Many recent research efforts have addressed the possibility of exploiting the wireless communication of IWMDs to compromise patients' privacy, or to send malicious commands that can cause unintended behavior. For example, Halperin *et al.* showed that the unencrypted wireless channel of a pacemaker can be exploited to compromise the confidentiality of data or to send unauthorized commands that cause the pacemaker to deliver therapy even when it was not needed [4]. Subsequently, a successful attack on an insulin pump, exploiting the wireless channel between the device and remote controller, was shown in [5]. By reverse-engineering the customized radio communication and interpreting the unencrypted packets sent from a remote controller to an insulin pump, the attacker can launch radio attacks to inject insulin into the patient's body beyond the dosage regimen. Finally, attacks that drain the battery of IWMDs by sending packets that fail authentication have also been proposed [4].

In this article, we demonstrate that medical privacy concerns extend far beyond the wireless communication to/from IWMDs. We make two main contributions. First, we describe the possibility of privacy attacks that target *physiological information leakage*, i.e., signals that are continuously emanating from the human body due to the normal functioning of its organs. These attacks are a concern even when there is no medical device present, and hence have a much wider scope.

As our second contribution, we target IWMDs. We propose several novel attacks on privacy by leveraging information leaked from them during their normal operation. We demonstrate attacks on two medical devices based on acoustic and electromagnetic (EM) leakage from them. Moreover, we investigate a novel metadata-based attack that extracts critical health-related information by monitoring the communication channel, although the data may be completely encrypted. We note that the proposed attacks are applicable even when medical devices have no wireless communication, or when the wireless communication is encrypted, unlike previous attacks that compromise unencrypted wireless channels [4], [5].

The rest of the article is organized as follows. Section II describes the threat model. Section III discusses the sources

Acknowledgments: This work was supported by NSF under Grant no. CNS-1219570.

Arsalan Mohsen Nia is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA (e-mail: arsalan@princeton.edu).

Susmita Sur-Kolay is with the Advanced Computing and Microelectronics Unit, Indian Statistical Institute, Kolkata 700108, India (e-mail: ssk@isical.ac.in).

Anand Raghunathan is with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA (e-mail: raghunathan@purdue.edu).

Niraj K. Jha is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA (e-mail: jha@princeton.edu).

and various types of physiological information leakage. Section IV presents our bevy of proposed privacy attacks. Section V suggests some countermeasures against the attacks, and Section VI concludes the paper.

II. Threat Model

In this section, we first describe potential adversaries. Then, we describe potential risks that may arise from loss of privacy.

A. Adversary

We consider an adversary to be any potentially untrusted person who has a short-term physical proximity to the patient. The proposed attacks, while not impossible, may be difficult to deploy in a secure location such as the patient's home or a medical facility such as a hospital. However, none of our attacks require access to specialized medical equipment such as the ones used in hospitals. We also assume that long-term physical access to the patient or monitoring of the patient, e.g., using a camera that continuously monitors the subject's activities, is not feasible. In our attack scenarios, the adversary gains the required physical access to the patient in any public location. Crowded places, such as train stations, bus stops, and shopping malls, may provide opportunities for the adversary to come closer to the subject, while hiding the required equipment. A potential adversary might be an employer who intends to discriminate against a chronically-ill patient, a private investigator who has been hired to spy on the subject, a political operative who wants to expose the medical condition of the subject for political advantage or a criminal group seeking to sell valuable medical information to the highest bidder [6].

B. Potential Risks

The patient's physiological signals may be exploited in various ways. We describe some of the consequences of such information leakage next.

- **Job/insurance loss:** Revelation of medical conditions may negatively impact a person's employment prospects or make it more difficult for him to obtain insurance. Leakage of this sensitive information from the human body or IWMDs, such as the presence of a serious illness, level of the illness, exposure of a condition that may carry social stigma, and exposure of physical, emotional or mental conditions would naturally raise serious privacy concerns.
- **Unauthorized interviews:** An unauthorized interviewer may be able to combine lie detection (also called deception detection) questioning methods with the privacy attack techniques proposed in this work to ascertain the truth or falsehood of responses given by the subject, without his consent. Several researchers have investigated variations in vital health signals, such as the respiratory rate and heart rate, in the presence of acute emotional stress (e.g., when the person is lying) or a mental problem [7]–[9]. For instance, Sung et al. have demonstrated changes in the heart and respiratory rates in live poker game scenarios [10].
- **Indirect consequences:** Although disclosure of medical information using the proposed privacy attacks might not be directly lethal, unlike attacks on the integrity of the medical device [4], [5], it may lead to a subsequent tailored integrity attack. For instance, as described later, extracting medical device information, model, type, and configuration using EM leakage from the device may provide enough information to an adversary to design a more effective integrity attack using the extracted parameters. Moreover, detection of usage of certain medical devices by adversaries may endanger the safety of the patient, e.g., if the device is very expensive and attracts theft, or embarrass the subject if the medical condition carries a social stigma [6].

III. Information leakage

In this section, we first discuss the possible sources of information leakage, followed by brief descriptions of different types of signal leakages addressed in this paper.

A. Leakage sources

In this work, we consider two sources of information leakage: (i) human body and (ii) IWMDs. Each source continuously leaks information through different types of signals.

Several organs in our body generate biomedical signals. Some of these signals can be remotely captured and analyzed. For example, our lungs generate an acoustic wave called *respiration sound*, which can be captured by a microphone.

In addition to body organs, IWMDs may also reveal critical health information under normal operation even when not using any wireless communication to transmit data. For example, the electrical motor of an insulin pump generates an acoustic signal when injecting insulin. As described later in Section IV, performing simple signal processing on this acoustic signal can reveal the prescribed insulin dose.

B. Leakage types

In general, leaked physiological signals can be divided into two types: (i) acoustic and (ii) EM signals. Fig. 1 demonstrates the sources of leakage, as well as the different types of signals that we consider in this work. Body organs, such as heart and lungs, produce an acoustic signal that can be captured remotely and analyzed. IWMDs, such as an insulin pump or BP monitor, may also generate acoustic and EM signals during their normal operation even if they are not transmitting any data. The following subsections describe these signals in detail.

B.1 Body-related information

The human body consists of several continuously-operating organs. Various acoustic and EM signals are generated as

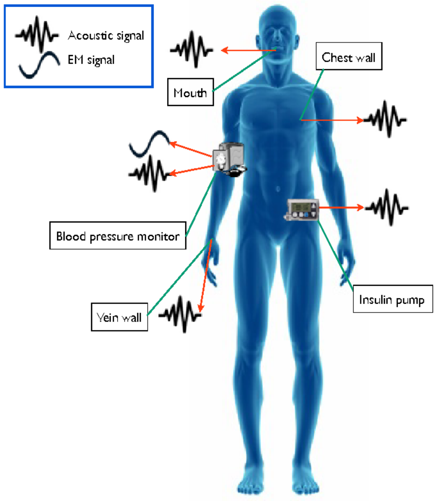

Fig. 1. Sources of leakage and different types of signals that are continuously leaking from the human body and IWMDs.

a result. The majority of these signals are too weak to be captured without physical contact or may be absorbed by other organs before emanating from the body. For example, electrical activities originating from nerves carry real-time information about health status. The two commonly-used methods for measuring these signals are electroencephalography (EEG) and electrocardiography (ECG). The amplitudes of EEG and ECG signals vary from tens of microvolts to few millivolts. The frequencies of most of these signals are below 40 Hz [11]. Another example is the acoustic signal generated by blood circulating through internal organs. However, this is absorbed by the surrounding muscles and tissues.

If an EM or acoustic signal generated by an organ emanates from the human body, it may be captured and analyzed to reveal health-related information. For example, one such signal is the respiration sound that is generated by chest vibration and airflow through the mouth. In the following subsections, we discuss different types of signals that might leak from the human body during normal operation.

B.1.1 Acoustic signals emanating from the human body

Some of the body organs generate acoustic signals during their normal operation. In this work, we examine the feasibility of capturing such naturally-occurring acoustic signals from a distance to reveal confidential health information of a person. Specifically, we show how capturing acoustic signals generated by two organs, heart and lungs, can reveal critical information.

As discussed later in Section IV, a simple signal processing algorithm enables us to count the number of peaks in the raw heart sound signal and thus compute the heart rate. The heart rate may be an indicator of several critical illnesses or a sudden emotional stress. For example, when a person lies, his heart rate gets elevated above the normal [12]. Therefore, if an adversary can monitor the heart rate remotely, he may even be able to assess whether the person is telling the truth.

Respiratory sounds also reveal valuable information about the health condition of an individual. The process of recording respiratory sounds and analyzing them is referred to as computerized respiratory sound analysis [13]; it provides crucial information on respiratory dysfunction, and changes in the respiratory characteristics (e.g., duration, timing).

B.1.2 EM signals emanating from the human body

The human body continuously emits infrared radiation that carries health information. These raw data can be captured and processed by an attacker at a distance. The existence of such a natural continuous leakage of information may allow an attacker to acquire critical information about the patient's health condition even if all IWMDs and their communications are completely secure, e.g., using encryption.

The use of thermal images has increased dramatically in the medical applications during the last decade. Thermal imaging cameras highlight warm objects against cooler backgrounds. As a result, the human body is easily visible in the environment. Moreover, some physiological variations in the human body can also be detected with thermal imaging techniques employed in medical diagnostic procedures. Several research projects on thermal imaging have been discussed in the medical literature. Using these methods [14]–[16], an eavesdropper can easily reveal the health status of a person. For example, in [14], Arora et al. showed the effectiveness of detecting breast cancer using digital infrared thermal imaging. The possibility of mass fever detection using thermal imaging techniques is described in [15].

B.2 IWMD-related information

As mentioned earlier, IWMDs are used for monitoring and therapeutic purposes. An IWMD may leak health-related data or metadata that compromise the patient's privacy. Next, we describe how IWMDs can leak information through acoustic and EM signals.

B.2.1 Acoustic signals emanating from IWMDs

First, we describe acoustic leakage from IWMDs. Acoustic waves propagate through a transmission medium using adiabatic compression and decompression. These waves are generated by a source. The source vibrates the medium, leading to propagation of vibrations from the source.

Electronic devices with microelectromechanical parts generate unintentional acoustic signals during normal operation. Some recent research efforts have demonstrated the feasibility of revealing critical information by interpreting acoustic emanations from peripheral computer devices. For example, researchers have shown that acoustic emanations from matrix printers carry substantial information about the printed text [17]. Moreover, Zhuang et al. have demonstrated the feasibility of recovering keystrokes typed on a keyboard from a sound recording of the user typing.

In this work, we demonstrate how acoustic signals generated by an IWMD (e.g., an insulin pump) may carry significant information about the patient's health status and the functioning of the medical device.

B.2.2 EM signals emanating from IWMDs

Next, we discuss EM radiations from IWMDs. We divide the EM radiations into two classes: (i) unintentionally-generated and (ii) intentionally-generated. Generally, an electronic equipment may emit unintentional EM signals that can be used as side-channel information, allowing eavesdroppers to reconstruct processed data at a distance [18]. This has been a concern in the design of military hardware for over half a century [19]. IWMDs can also unintentionally generate EM signals while performing their regular tasks. These signals may reveal critical information about the status of the medical device and patient's health condition. In this work, we demonstrate how an insulin pump can leak information about its function by emitting unintentional EM radiations.

In addition to unintentional EM radiations, medical devices may use EM signals intentionally to wirelessly transmit medical data. Eavesdropping on unencrypted wireless communication has been addressed in several research articles [4], [5]. In this work, we focus on the metadata that leaks through wireless communication even when the packets are encrypted.

IV. Privacy attacks

In this section, we propose and discuss various attacks on the privacy of medical data based on the information leaked from the human body and IWMDs. Table I summarizes the sources of leakage, the types of signals, and the information extracted from each attack, that are investigated in this work. For each attack, we first describe one or two methods to capture the signals and then our processing algorithms to interpret the captured signals in order to reveal the patient's health condition.

TABLE I
Sources of leakage, types of signals, and information extracted

Source	Type of signal	Information extracted
Human Body	Acoustic	Respiration/Heart rate
IWMDs	Acoustic	Insulin dose and BP
	EM (unintentional)	BP
	EM (wireless)	Device info. and insulin dose

A. Acoustic signal based body-related attacks

Next, we target the acoustic signals leaked during the normal functioning of lungs and heart. We first describe two methods for remotely capturing the sounds from these organs. Then, we demonstrate how we can accurately extract respiration and heart rates from the captured signals.

A.1 Capturing acoustic signals emanating from the lungs and heart

Method 1: We have used a displacement-based laser microphone that uses a laser beam to detect sound vibrations from a distance. Laser microphones were invented to eavesdrop on a conversation with a minimal chance of exposure. Although they have been used for surveillance purposes for a long time [20], for the first time, we employ these microphones in the context of a privacy attack on patients' medical data. We have built an inexpensive laser microphone to detect vibrations emanating from the human heart and lungs. This device is based on detecting the varying amounts of reflected laser beam received by a single ambient light sensor. As illustrated in Fig. 2, the laser beam forms a small incident angle with the surface. Surface vibration along the normal vector causes displacement of the reflected beam, and as a result, the amount of laser signal reaching the receiver varies for different displacements. Fig. 3 shows the receiver for capturing an acoustic signal using a laser microphone. It connects to the processing unit using an aux cable.

Method 2: The second capture method we propose is based on a parabolic microphone (KJB-Det [21]) to capture weak acoustic signals generated by the lungs. It uses a parabolic reflector to collect and focus sound waves onto a receiver. It amplifies the acoustic signal by concentrating all of the sonic energy at the focal point, thus increasing the signal-to-noise ratio (SNR). KJB-Det comes with a 20-inch parabolic dish. In addition, electronic amplifiers used in KJB-Det can increase the overall level of both noise and acoustic signal, without degrading the SNR.

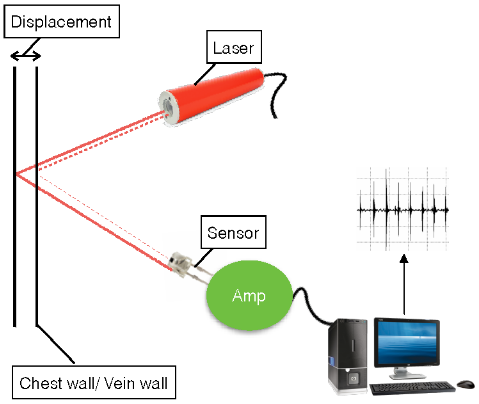

Fig. 2. Schematic for displacement-based laser microphone: the laser beam forms a small incident angle with the surface. The fraction of light beam received by the light sensor depends on the vibration of the surface.

A.2 Extracting respiration and heart rates

Next, we first describe a method to extract the heart and respiration rates from the captured acoustic signal. Then, we discuss the parameters that affect the accuracy and detection range of each of the two capture methods described above.

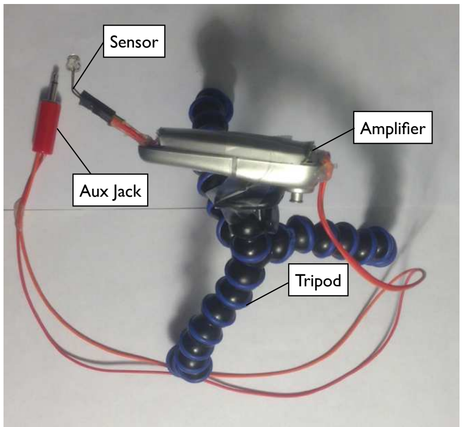

Fig. 3. Receiver set-up used for the displacement-based laser microphone.

For obtaining the heart and respiration rates, we use a simple algorithm to find the local maxima. In order to reduce the effect of noise, the algorithm ensures that the distance between two consecutive peaks is more than the value of a parameter called $distanceThreshold$. The maximum possible human heart rate (200 pulse per minute) and respiration rate (80 breaths per minute) are used to define $distanceThreshold$. Thus, $distanceThreshold$ is set to $5ms$ and $12.5ms$ for the heart and respiration rates, respectively.

In Method 1, we use the laser-displacement microphone for capturing acoustic signals from both the lungs and heart. The sound quality obtained by this microphone depends on two factors: (a) reflection fraction, which is the fraction of the incident beam that is reflected by the surface, and (b) the displacement of the received beam. The first parameter depends on the nature of the surface. For example, the human skin absorbs a large fraction of the incident beam; therefore, the sensor should be placed close to the skin to receive the beam. However, the displacement of the received beam on the sensor decreases as the sensor gets nearer the reflecting surface. We were able to accurately extract the respiration rate from 5 cm away. If the person wears a metallic/reflecting necklace, we can point the incident beam towards the necklace instead, which is a better reflector than the human skin. We were able to accurately extract the respiration rate from 6 m away when the person wore a flat steel necklace. We also used a displacement-based laser microphone to detect the heart rate. In the absence of an attached reflector surface, the acoustic signal was used by the laser microphone to detect the heart rate with over 95% accuracy at a distance of 5 cm. At greater distances, the amount of received beam reduces drastically and the accuracy drops.

The audio gain of a parabolic microphone increases as the frequency increases. The gain of an ideal 20-inch dish with a perfect parabolic shape and focus is characterized by a curve starting from 0 dB at 200 Hz. In order to enhance the amplification of our parabolic microphone, we replaced its dish with a larger 1 m dish that provides a 6 dB amplification at 200 Hz. At lower frequencies, the most important parameters are dish size and the quality of the microphone. The modified parabolic microphone was able to detect the respiration rate at a distance of 5 m. However, the parabolic microphone was unable to detect the heart rate.

B. Acoustic signal based IWMD-related attacks

Acoustic signals generated unintentionally by an IWMD can provide valuable information to an unauthorized party. Each IWMD consists of different components. Some of these components (e.g., electrical motors and relays) can produce a capturable sound during normal operation. An unintentionally-generated acoustic signal can be used as a side-channel information to reveal the status of the medical device and the patient's condition. In addition to this class of acoustic signals, some IWMDs intentionally produce acoustic signals to notify the users of conditions that require immediate attention. Many medical devices have alarm systems; among them are insulin pumps, pulse oximetry devices, and BP monitors. These alarms offer necessary warnings to inform patients of changes in their health condition. They usually provide sophisticated mechanisms for safety checks. These alarms make the patient aware of an unusual situation. Generally, the audible frequency range for a human is between 20 Hz and 20 kHz. Frequency ranges of 2 kHz to 4 kHz are most easily heard. For this reason, most alarms emit sound in this frequency range.

Several sound-recording equipments are available on the market, ranging from simple microphones to sophisticated parabolic microphones. In the following subsections, we describe two different attacks using acoustic signals. In the first attack, we capture and amplify the sound of an electrical motor using a parabolic microphone. In the second attack, we use a simple microphone to record the required acoustic signal. We were able to accurately determine the amount of injected insulin from 1 m and 10 m away for the first and second attacks, respectively. Using a more powerful microphone or amplifier will obviously increase this range.

B.1 Acoustic leakage from an insulin delivery system

We now describe how acoustic signals leaking from an insulin delivery system can reveal the patient's health condition. Fig. 4 shows a schematic view of an insulin pump. The display screen allows the user to set the value of different device parameters. The controller controls the motor, which pushes the piston rod forward to release a prescribed amount of insulin. In this medical device, the electrical motor unintentionally generates acoustic signals and the speaker intentionally produces different alarms as reminders for calibration and high/low glucose, and as predictive high/low glucose alerts. The components marked in red (motor and buzzer) generate the acoustic signals that we can interpret to reveal the medical data.

Here, we present two attacks on an insulin pump using these acoustic signals. First, we demonstrate how capturing and interpreting the unintentional acoustic signal generated by its electrical motor can reveal the patient's prescribed dosage, and hence the level of diabetes. Second, we use the acoustic signals generated by its safety system to remotely examine the status of the device and extract the prescribed dosage.

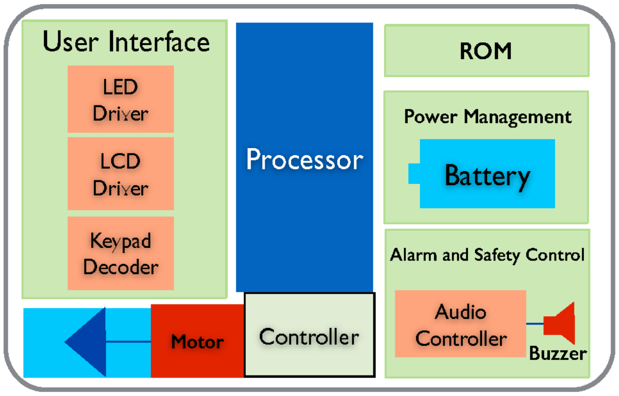

Fig. 4. A schematic view of an insulin pump

B.1.1 Extracting information from motor sounds of an insulin pump

We show below how processing the weak acoustic signal generated by the electrical motor of an insulin pump can reveal the exact amount of injected insulin, and as a result, provide an estimation of initial blood sugar, and level of diabetes. We propose two signal processing algorithms for this purpose.

In an insulin pump, a step motor is used in the injection procedure. Our experiments demonstrate that there is a linear relationship between the number of rotation steps of the electrical motor and the amount of injected insulin (Fig. 5). Fig. 6 illustrates the acoustic signal generated by the electrical motor while injecting 0.8 unit of insulin. Each peak corresponds to one step of the motor. The first processing algorithm finds the number of peaks. Thus, for this case, the total number of steps is calculated as 16, thereby inferring that 0.8 unit of insulin was injected.

Algorithm InsPumpI shows the pseudo-code for our first proposed algorithm. It calculates the exact amount of injected insulin based on the number of motor steps. Its four inputs are: (i) $acousticSignal$, which is the acoustic signal of the electrical motor sampled at 22 KHz, (ii) $distance$, which indicates the minimum acceptable distance between two consecutive peaks (steps), (iii) $threshold$, which is the minimum acceptable amplitude of a peak, and (iv) $widthThreshold$, which is the width of a step in the absence of environmental noise. We obtain the number of steps from the number of peaks in $acousticSignal$ using subroutine $stepCount$. Then, using another subroutine $stepWidth$, we calculate the width of each step that is defined as the time when the peak and its neighboring points are greater than $widthThreshold$. After finding the number of peaks and the width of each peak, we estimate the number of steps that might be corrupted by comparing the width of each peak to $widthThreshold$. If the width is more than $widthThreshold$, it is likely to contain noise in the area around the peak. This algorithm is able to automatically detect the number of peaks in $acousticSignal$ that are corrupted. If the number of corrupted locations in $acousticSignal$ is more than three, there will not be enough information in $acousticSignal$ to reveal the exact insulin dose. Therefore, the attacker should discard that waveform, and try again later when background noise is less powerful.

In order to evaluate and compare our acoustic signal based algorithms, we constructed a test set consisting of 20 acoustic signals generated by the insulin pump when injecting four different doses of 0.2, 0.4, 0.6, and 0.8 unit of insulin (five injections for each dose). We captured the first 10 acoustic signals in a silent office (low-noise environment). We captured the other acoustic signals in the presence of background noise generated by a conversation. The algorithm could extract the injected dose exactly for the first 10 cases. In the presence of the conversation, the algorithm correctly detected the corrupted signal in four cases, and extracted the exact injected dose in the other six cases.

Algorithm InsPumpI. Calculating the exact amount of insulin dosage from the acoustic signal leaked by the insulin pump.

Given: $acousticSignal$, $distance$, $threshold$, and $widthThreshold$

```
$steps \leftarrow stepCount(acousticSignal, threshold, distance)$
$widths \leftarrow stepWidth(acousticSignal, threshold, distance)$
$for\ each\ width\ in\ widths$
    $if(width > widthThreshold)$
        $noisy \leftarrow noisy + (width/widthThreshold)$
    $end$
$end$
$if(noisy > 3)$
    $Print$ “ $Warning: Inaccurate$ ”
    $return\ -1$
$else$
    $dosage \leftarrow [steps/4] * 0.2$
    $Print\ dosage$
    $return\ 0$
$end$
```

Output: $dosage$
Return Status: 0 (accurate) or -1 (inaccurate)

In addition to counting the number of steps, we can calculate the duration of an injection. Calculating the duration is more robust against noise. In our second porposed algorithm, we show how an adversary can use an estimation of the injection duration to find the exact amount of injected insulin without knowing the exact number of steps.

The amount of injected insulin is quantized to a multiple of 0.2 unit of insulin. As a result, the injection duration is quantized and is a multiple of 7 seconds. For example, the

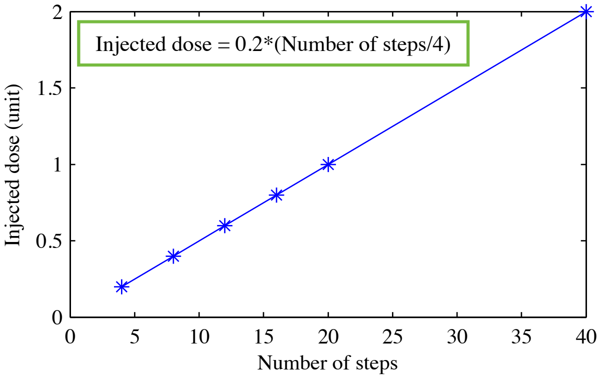

Fig. 5. Dose of injected insulin vs. the number of rotation steps of the electrical motor.

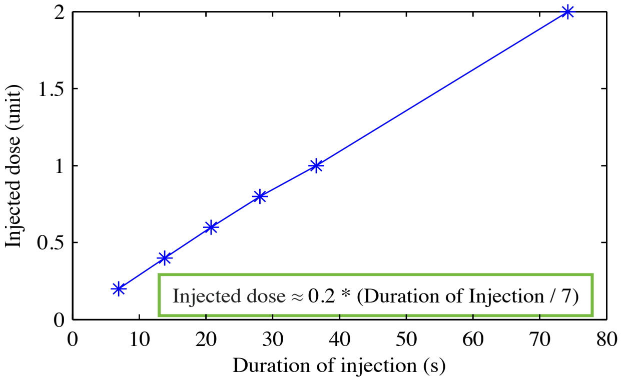

Fig. 7. Dose of injected insulin vs. injection duration.

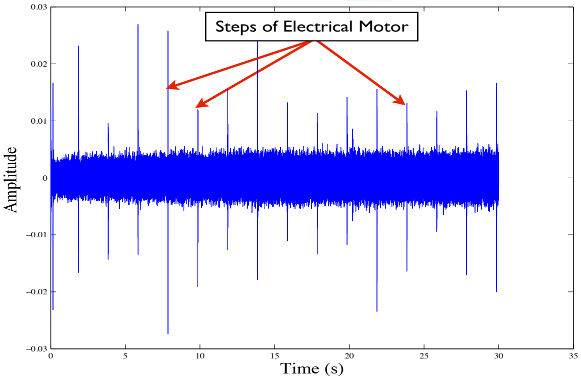

Fig. 6. Acoustic signal generated by the electrical motor of an insulin pump while injecting 0.8 unit of insulin.

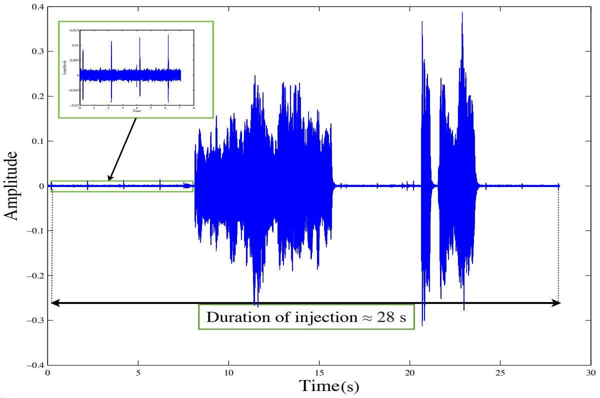

Fig. 8. Acoustic signal generated by the electrical motor of an insulin pump when 0.8 unit of insulin is injected. For a large fraction of time, the acoustic signal is dominated by background noise, and counting the number of rotation steps is not feasible.

injection of 0.2 and 0.4 unit of insulin takes about 7 and 14 seconds, respectively. Fig. 7 shows the amount of injected insulin with respect to injection duration. It shows there is an almost-linear relationship between the amount of injected insulin and injection duration. Therefore, if the attacker can only estimate the injection duration by calculating the time during which the sound of the electrical motor is present, he can find the exact amount of injected insulin even when a large fraction of the acoustic signal is dominated by background noise and counting the total number of steps is not feasible (Fig. 8). Using the test set described earlier, our duration-based algorithm was able to extract the exact amount of insulin in 18 of the 20 cases (10 under low-noise signals and 8 under noisy signals). Similar to the previous method, this algorithm was also able to automatically detect the situations in which the presence of noise affects the computed results.

In summary, capturing and processing the acoustic signal generated by the electrical motor of an insulin pump may reveal the injected dosage, and as a result, reveal the medical condition of the patient. The medical literature suggests that one unit of insulin is required per 50 mg/dl above 120 mg/dl of blood sugar [22]. Therefore, after measuring the insulin dosage, we can also estimate the level of blood sugar before injection.

B.1.2 Eavesdropping on alarms of an insulin pump

We describe below how the safety system of an insulin pump, which intentionally generates acoustic signals to inform patients, can unintentionally leak critical information about the health condition of a patient. As mentioned earlier, alarms are intended to alert patients of special events. The controller unit of the insulin pump (Fig. 4) is responsible for handling alerts and alarms, and the speaker generates audible signals in various situations, including blockage, low/high sugar level, initialization, and end of an injection.

Each injection procedure has four different phases: (i) initialization, (ii) confirmation, (iii) injection, and (iv) end of injection. Fig. 9 shows the acoustic signal generated by the alarm system of an insulin pump when a user tries to inject 0.8 unit of insulin. The four phases of the injection procedure are demonstrated in this figure. After the patient sends the injection command, the beginning of the initialization phase is reported by a single beep sound. Then, the user sets the dosage. In the confirmation phase, multiple beeps are generated based on the desired dosage. In this phase, one beep is generated by the safety system for every increment of 0.2 unit in insulin dose. However, the frequency of beeps in this phase is 2× higher than that in the initialization phase. Next, the injection begins and finally the end of injection is reported to the patient by a single beep.

We used three methods to find the exact amount of injected insulin by interpreting the acoustic signal: (i) initialization-based method that counts the number of peaks (beeps) in the initialization phase, (ii) confirmation-based method that counts the number of peaks in the confirmation phase, and (iii) duration-based method that calculates injection duration. Although the first two methods are straightforward, the third method is more accurate, especially in noisy environments. Similar to the previous attack (Section B.1.1), by extracting the quantized values of injection duration and dose, the exact prescribed dosage can be calculated, even if the attacker cannot count the number of beeps, but only estimate the injection duration. The injected dosage can be directly computed based on the almost-linear relationship between injection duration and injected dosage (Fig. 7).

In order to evaluate the accuracy of each algorithm, we constructed a similar test set to the one we used earlier. We captured the acoustic signal from 10 m away. The raw signal was amplified using a cheap amplifier before processing. All three methods could accurately extract the injected dose in the low-noise environment. Table II shows the number of correct and incorrect calculations and accuracy of each method in the noisy environment. The duration-based method showed the best accuracy, where accuracy is defined as the percentage of correctly-calculated doses.

TABLE II
ACCURACY OF THE THREE METHODS FOR EAVESDROPPING ON THE ALARM SYSTEM OF AN INSULIN PUMP

Method	Correct	Incorrect	Accuracy (%)
Initialization-based	6	4	60
Confirmation-based	7	3	70
Duration-based	10	0	100

In addition to compromising health-related information of a patient, the status of the medical device, such as blockage and low-battery state, can also be directly extracted by capturing and analyzing the alarms generated by the insulin pump.

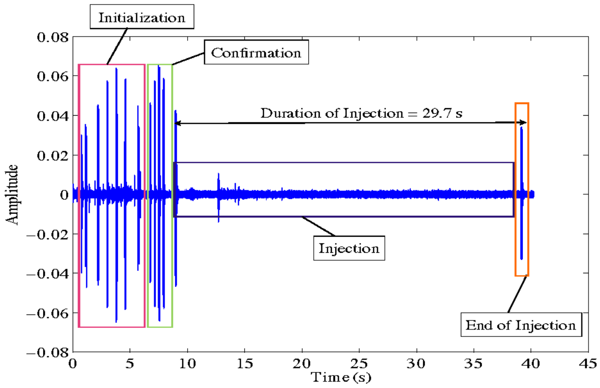

Fig. 9. Acoustic signal generated by the safety system of an insulin pump when the user tries to inject 0.8 unit of insulin.

B.2 Interpreting the leaked acoustic signal of an ambulatory BP monitoring device

In this section, we target an ambulatory BP monitoring device. Fig. 10 shows a block diagram of such a device. The components shown in red are the major sources of acoustic leakage. We interpret the sound generated by its electrical pump to estimate the patient's BP, which is the pressure generated by circulating blood upon the walls of blood veins. BP is commonly represented by two numbers (systolic and diastolic), and is measured in millimeters of mercury (mm Hg). Non-invasive ambulatory BP monitoring is being increasingly used to continuously monitor patients' BP. A digital BP monitor has a cuff and digital pressure sensor. When a user inserts his arm in the cuff, it is automatically inflated by an electric motor. The digital monitor determines the BP and heart rate by measuring the small oscillations when the pressure is slowly released from the cuff. Common BP monitoring devices use a simple algorithm to derive an upper bound on systolic BP. They inflate the cuff to reach the upper bound in every measurement. However, in order to ensure patient's comfort, some new BP devices often use a technology known as fuzzy logic, which anticipates systolic BP to prevent over-inflation. In these devices, the highest pressure in the cuff is approximately 10 mm Hg to 15 mm Hg more than the actual systolic pressure. In this paper, we have targeted a commercially available BP monitoring system. We choose not to disclose its brand and model number. Next, we discuss how the sound generated by the electrical pump can provide enough information for an eavesdropper to accurately estimate the BP (both systolic and diastolic).

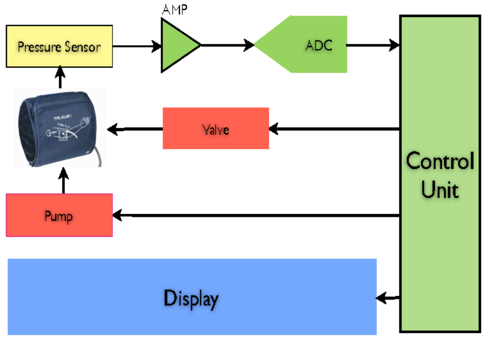

Fig. 10. Block diagram of an ambulatory BP monitoring device.

Our experiments demonstrate that each measurement consists of three consecutive phases: (i) inflation phase in which the cuff pressure increases to reach its upper bound value, (ii) step-wise deflation phase in which the monitoring device opens an air valve to slowly decrease the cuff pressure and measure the BP, and (iii) restart phase. Fig. 11 shows the acoustic signal generated during the measurement.

In the BP monitoring device used in our experiments, the cuff pressure decreases about 9 mm Hg for each step of deflation in the second phase of measurement. In addition,

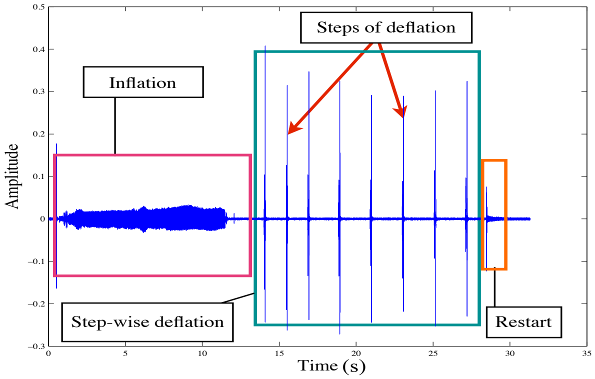

Fig. 11. Acoustic signal generated by the ambulatory BP monitoring device. Three phases of measurement are shown.

we found that the systolic BP was detected after three or four steps in the step-wise deflation phase, which suggests that the systolic BP should be in the range of $(P_h - 27)$ mm Hg to $(P_h - 36)$ mm Hg, where P_h is the maximum cuff pressure in the inflation phase. Moreover, based on our experimental results, the diastolic pressure is usually detected in the range of P_l mm Hg to $(P_l + 9)$ mm Hg, where P_l is the minimum cuff pressure during the step-wise deflation phase before the device enters the restart phase. In order to examine the accuracy of the above claim, we used 25 BP measurements. The systolic BP was in the range of $(P_h - 27)$ mm Hg to $(P_h - 36)$ mm Hg for 21 out of 25 measurements. Moreover, for 23 out of 25 measurements, the diastolic BP was in the range of P_l mm Hg to $(P_l + 9)$ mm Hg. Therefore, if we can develop a method to detect P_h and P_l, the systolic and diastolic pressure can be estimated as $(P_h - 27 + P_h - 36)/2$ mm Hg and $(P_l + P_l + 9)/2$ mm Hg, respectively.

Next, we describe how we can use the acoustic signal generated by the electrical pump to extract P_h and P_l and thus estimate the BP. The cuff pressure reaches its maximum value at the end of the inflation phase. In order to find the maximum value for an arbitrary measurement, we construct a look-up table that maps the maximum pressure (P_h) to the duration of inflation (T_h), where P_h varies from 100 mm Hg to 180 mm Hg. For each measurement, we first calculate T_h by finding the part of the acoustic signal in which the pumping sound is present. Then, we use the look-up table to find P_h. Thereafter, we count the number of steps before deflation. Then, we calculate the range of systolic and diastolic pressures, and report the middle points of these ranges as their estimate. Our experimental results show that for 19 out of 25 arbitrary measurements, this algorithm calculates both systolic and diastolic pressures with absolute error less than 8%, where error is defined as the difference between the estimated and actual values divided by the actual value. The main reason for the six failed cases was re-inflation. Re-inflation occurs when the patient suddenly changes his arm position during the step-wise deflation phase. In this case, the monitoring device increases the cuff pressure again. It is easy to modify the method to detect the situation in which re-inflation occurs. *Algorithm AmbBP* gives the pseudo-code for the improved version of our algorithm. We define a function, called $infCount$, which finds the number of inflation phases separated by deflation steps. This improved algorithm automatically detects whether the algorithm is unable to calculate the BP accurately.

Algorithm AmbBP. Estimating systolic and diastolic BP by processing the acoustic signal from an ambulatory BP monitor

Given: $acousticSignal, table$ where $table : T_h \rightarrow P_h$

```
$infNumber \leftarrow infCount(acousticSignal)$
$if(InfNumber > 1)$
        $Print$ “ $Warning$: $Inaccurate$ ”
        $return - 1$
$end$
$T_h \leftarrow calculateTimeOfInflation(acousticSignal)$
$P_h \leftarrow lookUp(T_h, table)$
$Steps \leftarrow CountPeaks(acousticSignal)$
$upperSystolic \leftarrow P_h - 27$
$lowerSystolic \leftarrow P_h - 36$
$P_l \leftarrow P_h - numberOfSteps * 9$
$upperDiastolic \leftarrow P_l$
$lowerDiastolic \leftarrow P_l + 9$
$systolic \leftarrow \frac{upperSystolic + lowerSystolic}{2}$
$diastolic \leftarrow \frac{upperDiastolic + lowerDiastolic}{2}$
$Print\ diastolic,\ systolic$
$return\ 0$
```

Output: *diastolic*, *systolic*, or the warning message
Return status: 0 (accurate) or -1 (inaccurate)

C. EM radiation based IWMD-related attacks

We target two classes of EM radiations: (i) unintentional EM radiations that are signals generated by different components of an IWMD (e.g., processor, controller), and (ii) intentional EM radiations that are encrypted wireless communications that transmit medical data. Next, we discuss two EM radiation based attacks using each class of EM radiations, namely from the pump in a BP monitor, and based on the metadata of wireless communications of an insulin pump.

C.1 Estimating BP from unintentional EM radiations

Next, we discuss an attack based on capturing and analyzing the EM radiation that is unintentionally generated by the BP monitoring device.

C.1.1 Capturing unintentional EM signals

We use an oscilloscope (MSO/DPO5000) to detect the EM signals. The EM side-channel information that we capture is available during the normal operation of the medical device even when the device is not transmitting any data (e.g., using a USB cable or wireless communication). We capture the raw EM signal directly from the antenna that is connected to the oscilloscope, instead of a filtered and demodulated signal with limited bandwidth. We use an antenna (75 Ohms VHA 9103 Dipol Balun) to improve the SNR for signals in the 25 MHz to 500 MHz frequency band. Moreover, we check if these EM signals can be captured using a small portable antenna, such as a simple loop of 0.5-meter copper wire.

EM signals may remain undetected using standard techniques. Spectral analyzers need significantly static carrier signals. The demodulation process may eliminate the interesting components of unintentionally-emitted EM signals. In addition, the scanning process of wide-band receivers may take a lot of time [23].

C.1.2 Processing the captured EM signals

Using the EM signals captured from the BP monitoring device, we were able to estimate the patient's BP. EM radiations reveal the activity of the electrical pump in the different phases of a measurement (inflation, step-wise deflation, and restart). The duration of the inflation phase can be revealed by calculating the time when the electrical motor produces the EM radiations, and as a result, the systolic BP can be extracted by using the method discussed earlier for extracting systolic BP from acoustic signals. Moreover, by monitoring the activity of the device in the deflation phase, the number of deflation steps could be detected. Estimating the BP using EM signals was as accurate as when it was estimated from acoustic signals. However, this method can be easily used in a crowded environment, where the acoustic signal may be dominated by background noise. The activity of the electrical pump in the inflation phase was completely detectable from 15 cm away when we used the VHA antenna. Moreover, when we replaced the VHA antenna with a 0.5 m wire, we were able to detect the activity from 10 cm away. The deflation steps were detectable using the VHA antenna and wire from 10 cm and 5 cm away, respectively.

C.2 Extracting insulin dosage regimen from the wireless communication metadata of the insulin pump

Next, we describe how capturing and processing the metadata leaked from the communication channel of an insulin pump can reveal critical medical information, including the injected dose of insulin, number of injections, and level of diabetes.

C.2.1 Capturing the metadata of wireless communication

In order to monitor fully-encrypted wireless communication and extract the metadata from the communication channel, we first need to find the frequency band of the transmission. If the model and type of the device are known, the frequency range can be extracted from manufacturer's documentation.

In general, an IWMD should make its existence and type unknown to unauthorized parties. If a device reveals its existence, its type should still remain hidden to unauthorized persons. This may be for many different reasons. For example, the device might be extremely expensive. More importantly, knowing the specific model of a device may provide critical information to potential adversaries. As we elaborate later, if the type, characteristics, and settings of an IWMD are known, designing a tailored attack becomes much easier. A tailored attack is a smart attack based on the specific features and configurations of a known device. Therefore, we assume that the model and type of the IWMD are not known to the attacker.

A fast approach for detecting the frequency band of a wireless transmission is through an oscilloscope that uses a loop of wire as an antenna. The eavesdropper can scan different frequency ranges when the communication channel is active and guess the frequency range. In addition, the frequency band of communication for an unknown IWMD can usually be obtained by scanning some specific bands based on the fact that FDA regulations impose specific limits on the frequency bands of medical devices. The majority of medical devices communicate at 450 MHz, 600 MHz, 900 MHz, 1.4 GHz, and 2.4 GHz.

After finding the frequency band of transmission, the encrypted packets can be captured using one or multiple universal software radio peripherals (USRPs) [24]. Next, we demonstrate how examining the frequency band of the channel and characteristics of the packets can reveal critical health information.

C.2.2 Processing the captured EM signals

Different manufacturers have different priorities and considerations. Thus, design priorities of IWMDs vary from one device to another. As a result, the metadata of the communication channel of one device are different from those of others. The metadata-based attack that we discuss next consists of two main steps: (i) the eavesdropper first extracts the metadata from the communication channel to reveal valuable information about the type and model of the IWMD, and (ii) when the device type is known, the attacker designs a tailored attack that specifically targets the known device. We discuss six classes of metadata leaked from the communication channel that can be used to find valuable information about the device: (i) frequency of communication, (ii) time between two consecutive transmissions, (iii) communication protocol, (iv) packet size, (v) detection range, and (vi) modulation protocol. However, in most cases, a subset of these classes can uniquely identify the model and type of the device.

We describe this attack using the insulin pump delivery system. For the insulin pump we used in this research, the frequency of communication (around 900 MHz), time between two packets (5 minutes), and modulation protocol (on-off keying) would be conclusive enough for an adversary to uniquely identify the insulin pump and its manufacturer. In addition, the detection range (20 m) and packet size (80 bits) match the information given in the documentation of the device.

Next, we describe a tailored attack against a known insulin pump. We assume all communications are fully encrypted. In the first step of the metadata-based attack, we find the model and type of the insulin pump. This specific model comes with a remote control. The remote control is a device that controls and programs the insulin pump and allows the user to deliver a discrete bolus dose or stop/resume insulin delivery. Each button on the remote control sends a specific command to the insulin pump. The size of remote control is usually small to assure patient's convenience, and as a result, there are only a few buttons on the remote control. Different sequences of buttons on the insulin pump are to be pressed in different situations. For the insulin pump that we monitored in our experiment, the patient should use at least three button presses to start the injection: (i) the first button tells the device to initialize the injection, (ii) the second button is used to set the dosage of injection, and (iii) the third button confirms the injection. The patient can press the second button multiple times to increase the dosage. In this scenario, interpreting the number of consecutive packets can uniquely reveal the occurrence of the injection, and the insulin dosage. For example, if seven packets are captured by the USRP in this case, the first and last packets would represent initialization and confirmation. Thus, the other five packets can be assumed to be sent to increase the amount of injected insulin. Therefore, monitoring the transmission channel, even when it is fully encrypted and packets do not carry any meaning to the attacker, can reveal the prescribed dosage of insulin. Moreover, the number of injections can be extracted by counting the number of transmissions that include more than three packets.

V. Possible Countermeasures

In this section, we briefly discuss some possible countermeasures to protect the patient against the privacy attacks described in this article. We hope these initial suggestions would spur further research on countermeasures against such attacks. We discuss different countermeasures for each source of leaked signals (human body and IWMDs).

Hiding information that leaks from the body is difficult because there are many local sources of leakage, e.g., lungs, heart, and skin. We can hide some of this information using cloth as a shield. However, since it is typically not possible to cover the whole body, medical information may at least leak from the face. For example, the EM radiation from the face leaks enough information to detect if a person has fever. Moreover, many components inside medical devices may generate acoustic or EM signals: the motherboard, communication cables, processor, and actuators. The simplest solution for eliminating the leakage of compromising information from IWMDs is use of a shield. However, incorporating a shield will increase the IWMD price and thus may not be desirable from a cost perspective. Another solution could be to analyze the local sources of leakage (e.g., motherboard, wires, and display board) during the manufacturing process, and add extra components to generate noise with specific characteristics so as to hide the information leakage. This approach also increases the cost of manufacture. Moreover, adding a noise generator may increase the energy consumption of the device and thus reduce its battery lifetime. Such masking techniques have been explored in the context of traditional side-channel attacks on cryptographic systems.

VI. Conclusion

In this paper, we discussed two sources, namely the human body and IWMDs, that continuously leak health information under normal operation. We targeted two types of signals for each source: acoustic and EM. We then described a variety of attacks on the privacy of health data by capturing and processing unintentionally-generated leaked signals. Moreover, we discussed the feasibility of using intentionally-generated acoustic signals (as a side-channel information) and EM signals (as a carrier of metadata) to compromise the patient's health privacy. Finally, we suggested some countermeasures.

References

[1] M. Zhang, A. Raghunathan, and N. K. Jha, "Trustworthiness of medical devices and body area networks," *Proc. IEEE*, vol. 102, no. 8, pp. 1174–1188, Aug. 2014.

[2] J. C. Haartsen, "The bluetooth radio system," *IEEE Personal Communications*, vol. 7, no. 1, pp. 28–36, 2000.

[3] P. Baronti, P. Pillai, V. W. Chook, S. Chessa, A. Gotta, and Y. F. Hu, "Wireless sensor networks: A survey on the state of the art and the 802.15.4 and ZigBee standards," *Computer Communications*, vol. 30, no. 7, pp. 1655–1695, 2007.

[4] D. Halperin, T. S. Heydt-Benjamin, B. Ransford, S. S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno, and W. H. Maisel, "Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses," in *Proc. IEEE Symp. Security and Privacy*, 2008, pp. 129–142.

[5] C. Li, A. Raghunathan, and N. K. Jha, "Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system," in *Proc. IEEE Int. Conf. e-Health Networking Applications and Services*, 2011, pp. 150–156.

[6] D. Halperin, T. Kohno, T. S. Heydt-Benjamin, K. Fu, and W. H. Maisel, "Security and privacy for implantable medical devices," *IEEE Pervasive Computing*, vol. 7, no. 1, pp. 30–39, 2008.

[7] F. Mokhayeri, M. R. Akbarzadeh-T, and S. Toosizadeh, "Mental stress detection using physiological signals based on soft computing techniques," in *Proc. 18th Iranian Conf. Biomedical Engineering (ICBME)*, Dec. 2011, pp. 232–237.

[8] B. Kaur, J. J. Durek, B. L. O'Kane, N. Tran, S. Moses, M. Luthra, and V. N. Ikonomidou, "Heart rate variability (HRV): An indicator of stress," in *Proc. SPIE Sensing Technology + Applications*, 2014, pp. 91 180V–91 180V8.

[9] G. N. Dikecligil and L. R. Mujica-Parodi, "Ambulatory and challenge-associated heart rate variability measures predict cardiac responses to real-world acute emotional stress," *Biological Psychiatry*, vol. 67, no. 12, pp. 1185–1190, 2010.

[10] M. Sung and A. Pentland, "Pokermetrics: Stress and lie detection through non-invasive physiological sensing," Ph.D. dissertation, Ph.D. thesis, MIT Media Laboratory, 2005.

[11] D. Svard, A. Cichocki, and A. Alvandpour, "Design and evaluation of a capacitively-coupled sensor readout circuit toward contact-less ECG and EEG," in *Proc. IEEE Biomedical Circuits and Systems Conference*, 2010, pp. 302–305.

[12] M. T. Bradley and M. P. Janisse, "Accuracy demonstrations, threat, and the detection of deception: Cardiovascular, electrodermal, and pupillary measures," *Psychophysiology*, vol. 18, no. 3, pp. 307–315, 1981.

[13] H. Pasterkamp, S. S. Kraman, and G. R. Wodicka, "Respiratory sounds: Advances beyond the stethoscope," *American J. Respiratory and Critical Care Medicine*, vol. 156, no. 3, pp. 974–987, 1997.

[14] N. Arora, D. Martins, D. Ruggerio, E. Tousimis, A. J. Swistel, M. P. Osborne, and R. M. Simmons, "Effectiveness of a noninvasive digital infrared thermal imaging system in the detection of breast cancer," *The American J. Surgery*, vol. 196, no. 4, pp. 523–526, 2008.

[15] J. B. Mercer and E. F. J. Ring, "Fever screening and infrared thermal imaging: Concerns and guidelines," *Thermology International*, vol. 19, no. 3, pp. 67–69, 2009.

[16] L. J. Jiang, E. Y. K. Ng, A. C. B. Yeo, S. Wu, F. Pan, W. Y. Yau, J. H. Chen, and Y. Yang, "A perspective on medical infrared imaging," *J. Medical Engineering and Technology*, vol. 29, no. 6, pp. 257–267, 2005.

[17] M. Backes, M. Dürmuth, S. Gerling, M. Pinkal, and C. Sporleder, "Acoustic side-channel attacks on printers," in *Proc. USENIX Security Symposium*, 2010, pp. 307–322.

[18] H. Tanaka, "Evaluation of information leakage via electromagnetic emanation and effectiveness of Tempest," *IEICE Trans. Information and Systems*, vol. 91, no. 5, pp. 1439–1446, 2008.

[19] ——, "Information leakage via electromagnetic emanations and evaluation of Tempest countermeasures," *Information Systems Security*, pp. 167–179, 2007.

[20] T. Wang, Z. Zhu, and A. Divakaran, "Long range audio and audio-visual event detection using a laser Doppler vibrometer," in *Proc. SPIE Defense, Security, and Sensing*, 2010, p. 77040J.

[21] "Detect ear - DET EAR," http://www.kjbsecurity.com/products/detail/detect-ear/117/, accessed: 02-1-2015.

[22] K. L. Herbst and I. B. Hirsch, "Insulin strategies for primary care providers," *Clinical Diabetes*, vol. 20, no. 1, pp. 11–17, 2002.

[23] L. Zhuang, F. Zhou, and J. D. Tygar, "Keyboard acoustic emanations revisited," 2005, pp. 373–382.

[24] M. Ettus, "Usrp users and developers guide," *Ettus Research LLC*, 2005.

Arsalan Mohsen Nia received his B.S. degree in Computer Engineering from Sharif University of Technology, Tehran, Iran, in 2012, and M.A. degree in Electrical Engineering from Princeton, NJ, in 2014. He is currently pursuing a Ph.D. degree in Electrical Engineering at Princeton University, NJ. His research interests include wireless sensor networks, Internet of things, computer security, distributed computing, mobile computing, and machine learning.

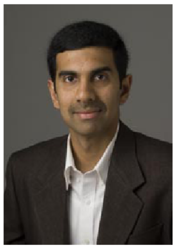

Anand Raghunathan is a Professor and Chair of VLSI in the School of Electrical and Computer Engineering at Purdue University, where he leads the Integrated Systems Laboratory. His research explores domain-specific architecture, system-on-chip design, embedded systems, and heterogeneous parallel computing. Previously, he was a Senior Research Staff Member at NEC Laboratories America and held the Gopalakrishnan Visiting Chair in the Department of Computer Science and Engineering at the Indian Institute of Technology, Madras. Prof. Raghunathan has co-authored a book (“High-level Power Analysis and Optimization”), eight book chapters, 21 U.S patents, and over 200 refereed journal and conference papers. His publications have been recognized with eight best paper awards and four best paper nominations. He received the Patent of the Year Award (recognizing the invention with the highest impact), and two Technology Commercialization Awards from NEC. He was chosen by MIT’s Technology Review among the TR35 (top 35 innovators under 35 years, across various disciplines of science and technology) in 2006, for his work on “making mobile secure”. Prof. Raghunathan has served on the technical program and organizing committees of several leading conferences and workshops. He has chaired the ACM/IEEE International Symposium on Low Power Electronics and Design, the ACM/IEEE International Conference on Compilers, Architecture, and Synthesis for Embedded Systems, the IEEE VLSI Test Symposium, and the IEEE International Conference on VLSI Design. He has served as Associate Editor of the IEEE Transactions on CAD, IEEE Transactions on VLSI Systems, ACM Transactions on Design Automation of Electronic Systems, IEEE Transactions on Mobile Computing, ACM Transactions on Embedded Computing Systems, IEEE Design Test of Computers, and the Journal of Low Power Electronics. He was a recipient of the IEEE Meritorious Service Award (2001) and Outstanding Service Award (2004). He is a Fellow of the IEEE, and Golden Core Member of the IEEE Computer Society. Prof. Raghunathan received the B. Tech. degree in Electrical and Electronics Engineering from the Indian Institute of Technology, Madras, and the M.A. and Ph.D. degrees in Electrical Engineering from Princeton University.

Susmita Sur-Kolay (SM05) received the B.Tech. degree in electronics and electrical communication engineering from Indian Institute of Technology, Kharagpur, India, and the Ph.D. degree in Computer Science and Engineering from Jadavpur University, Kolkata, India. She was in the Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA, from 1980 to 1984. She was a Post-Doctoral Fellow in the University of Nebraska-Lincoln, Nebraska-Lincoln, NE, USA, in 1992, a Reader in Jadavpur University from 1993 to 1999, a Visiting Faculty Member with Intel Corporation, Santa Clara, CA, USA, in 2002, and a Visiting Researcher at Princeton University in 2012. She is a Professor in the Advanced Computing and Microelectronics Unit, Indian Statistical Institute, Kolkata. She has co-edited two books, authored a book chapter in the Handbook of Algorithms for VLSI Physical Design Automation, and co-authored about 100 technical articles. Her current research interests include electronic design automation, hardware security, quantum computing, and graph algorithms. Prof. Sur-Kolay was a Distinguished Visitor of the IEEE Computer Society, India. She has been an Associate Editor of the IEEE Transactions on Very Large Scale Integration Systems, and is currently an Associate Editor of ACM Transactions on Embedded Computing Systems. She has served on the technical program committees of several leading conferences, and as the Program Chair of the 2005 International Conference on VLSI Design, the 2007 International Symposium on VLSI Design and Test, and the 2011 IEEE Computer Society Annual Symposium on VLSI. Among other awards, she was a recipient of the President of India Gold Medal from IIT Kharagpur.

Niraj K. Jha (S’85-M’85-SM’93-F’98) received his B.Tech. degree in Electronics and Electrical Communication Engineering from Indian Institute of Technology, Kharagpur, India in 1981, M.S. degree in Electrical Engineering from S.U.N.Y. at Stony Brook, NY in 1982, and Ph.D. degree in Electrical Engineering from University of Illinois at Urbana-Champaign, IL in 1985. He is a Professor of Electrical Engineering at Princeton University. He is a Fellow of IEEE and ACM. He received the Distinguished Alumnus Award from I.I.T., Kharagpur. He has served as the Editor-in-Chief of IEEE Transactions on VLSI Systems and an Associate Editor of IEEE Transactions on Circuits and Systems I and II, IEEE Transactions on VLSI Systems, IEEE Transactions on Computer-Aided Design, and Journal of Electronic Testing: Theory and Applications. He is currently serving as an Associate Editor of IEEE Transactions on Computers, Journal of Low Power Electronics and Journal of Nanotechnology. He has also served as the Program Chairman of the 1992 Workshop on Fault-Tolerant Parallel and Distributed Systems, the 2004 International Conference on Embedded and Ubiquitous Computing, and the 2010 International Conference on VLSI Design. He has served as the Director of the Center for Embedded System-on-a-chip Design funded by New Jersey Commission on Science and Technology. He is the recipient of the ATT Foundation Award and NEC Preceptorship Award for research excellence, NCR Award for teaching excellence, and Princeton University Graduate Mentoring Award. He has co-authored or co-edited five books titled Testing and Reliable Design of CMOS Circuits (Kluwer, 1990), High-Level Power Analysis and Optimization (Kluwer, 1998), Testing of Digital Systems (Cambridge University Press, 2003), Switching and Finite Automata Theory, 3rd edition (Cambridge University Press, 2009), and Nanoelectronic Circuit Design (Springer, 2010). He has also authored 15 book chapters. He has authored or co-authored more than 400 technical papers. He has coauthored 14 papers, which have won various awards. These include the Best Paper Award at ICCD93, FTCS97, ICVLSID98, DAC99, PDCS02, ICVLSID03, CODES06, ICCD09, and CLOUD10. A paper of his was selected for The Best of ICCAD: A collection of the best IEEE International Conference on Computer-Aided Design papers of the past 20 years, two papers by IEEE Micro Magazine as one of the top picks from the 2005 and 2007 Computer Architecture conferences, and two others as being among the most influential papers of the last 10 years at IEEE Design Automation and Test in Europe Conference. He has co-authored another six papers that have been nominated for best paper awards. He has received 14 U.S. patents. He has served on the program committees of more than 150 conferences and workshops. His research interests include FinFETs, low power hardware/software design, computer-aided design of integrated circuits and systems, digital system testing, quantum computing and secure computing. He has given several keynote speeches in the area of nanoelectronic design and test.