

Ghost Talk: Mitigating EMI Signal Injection Attacks against Analog Sensors

Denis Foo Kune^{*}, John Backes[†], Shane S. Clark[‡], Daniel Kramer, MD[§], Matthew Reynolds, MD[¶],
Kevin Fu^{*}, Yongdae Kim^{||}, and Wenyuan Xu^{**}

^{*}University of Michigan

[†]University of Minnesota, Twin Cities

[‡]University of Massachusetts Amherst

[§]Beth Israel Deaconess Medical Center, Harvard Medical School

[¶]Harvard Clinical Research Institute

^{||}Korea Advanced Institute of Science and Technology (KAIST)

^{**}University of South Carolina

Abstract—Electromagnetic interference (EMI) affects circuits by inducing voltages on conductors. Analog sensing of signals on the order of a few millivolts is particularly sensitive to interference. This work (1) measures the susceptibility of analog sensor systems to signal injection attacks by intentional, low-power emission of chosen electromagnetic waveforms, and (2) proposes defense mechanisms to reduce the risks.

Our experiments use specially crafted EMI at varying power and distance to measure susceptibility of sensors in implantable medical devices and consumer electronics. Results show that at distances of 1–2 m, consumer electronic devices containing microphones are vulnerable to the injection of bogus audio signals. Our measurements show that in free air, intentional EMI under 10 W can inhibit pacing and induce defibrillation shocks at distances up to 1–2 m on implantable cardiac electronic devices. However, with the sensing leads and medical devices immersed in a saline bath to better approximate the human body, the same experiment decreased to under 5 cm.

Our defenses range from prevention with simple analog shielding to detection with a signal contamination metric based on the root mean square of waveform amplitudes. Our contribution to securing cardiac devices includes a novel defense mechanism that probes for forged pacing pulses inconsistent with the refractory period of cardiac tissue.

Keywords—Attacks and defenses; embedded systems security; hardware security; analog sensors.

I. INTRODUCTION

Analog sensors have increasingly become an indispensable part of many modern systems, ranging from smartphones to medical devices to closed-loop control systems. The application layer running on these systems makes critical decisions, including actuation based on inputs from sensors including temperature, flow, position, electrocardiograms, electroencephalograms, and microphones. Unfortunately, analog sensors sensitive to electromagnetic interference (EMI) can provide an unchecked entry point into otherwise protected systems, allowing an attacker to manipulate sensor readings without changing the underlying physical

Lead student author: D. Foo Kune

Corresponding faculty authors: K. Fu, Y. Kim, W. Xu




Figure 1. Common device architectures implicitly trust sensed inputs. An attacker controlling sensed inputs can thus manipulate the application layer.

phenomena. The modified sensed data can appear directly at a device’s application layer, bypassing common security mechanisms (Figure 1) and giving the attacker some level of control over the system.

EMI affects circuits by inducing voltages on conductors — an effect also known as “back-door” coupling [1], where components become unintentional antennas capturing EMI radiation [2], [3]. EMI sources can be divided into *intentional* or *unintentional*, and *low-power* or *high-power*. There is abundant work devoted to unintentional or high-power EMI, but the effect of intentional low-power EMI for analog signal injection has yet to be explored. Unintentional high-power EMI sources, such as lightning strikes, electric trains, transformers [4], and sometimes communicating radios [5] are known to have an impact on modern circuits and analog sensors. Unintentional low-power leaks can allow eavesdropping on a system [6] and unintentional low power sources are also well known and accounted for in circuit designs [2]. In medicine unintentional high power EMI radiated from tools used for procedures like electrocautery [7] and from MRIs [8], [9] can affect cardiac implantable electrical devices (CIEDs). In addition, cell phones and other modern transmission devices have been investigated for their effect on CIEDs [10], [11], [12] and risks to patients.

Intentional EMI at high power can disable an adversary’s electronic components [13], [1], [14], [15]; intentional EMI

can also be used to inject faults into digital logic, leading to security violations [16], but this typically requires several volts of induced potential to succeed. Analog circuits, however, operating on the order of a few millivolts, can be vulnerable to interference at much lower energy levels and can allow the injection of forged sensor readings.

This work focuses on the signal injection with intentional low-power EMI on analog sensors. We analyze the root causes that enable signal injection on analog sensors, demonstrate that several commodity sensors are vulnerable to our EMI attacks, and investigate defense strategies.

Low-Power EMI Attack Analysis. Using back-door coupling, we design two types of EMI attacks:

1. *Baseband EMI attacks* inject signals within the same frequency band as sensor readings. Thus they are effective against analog sensors equipped with filters that attenuate signals outside intended frequency bands.
2. *Amplitude-modulated EMI attacks* modulate an attack signal on a carrier within the frequency band to which the victim’s analog sensors respond. Since the frequency of the EMI signal can match to the resonant frequency of a sensor, a successful attack requires a lower transmission power than baseband EMI attacks.

EMI Attack Validation. We demonstrated EMI attacks on medical devices monitoring electrograms and on commodity electronics using microphones. Specifically, despite proper filters in CIEDs, we successfully injected forged signals in leads in free air, causing pacing inhibition and defibrillation from 1 to 2 m away by transmitting at about 10 W and using a simple whip monopole antenna. With the device submerged in saline solution, the results decreased to under 5 cm. We also found that many commodity devices lack filters and are vulnerable to high frequency EMI signals. Using a transmitter with a power output of less than 100 mW, we were able to inject audio signals on microphones at a distance of up to 1 to 2 m. Our audio signals consisted of simple sinusoids, Dual-Tone Multiple Frequency (DTMF) signals commonly used in modern telephony, and arbitrary waveforms such as human speech and music.

Mitigation. While defenses against EMI attacks exist, we nonetheless found many devices vulnerable. We applied and measured the attenuation of our attack signal by known defenses including shielding, filtering and common mode noise rejection. Those techniques ameliorate but do not eliminate the injected EMI signals. Thus, we propose software-based defenses that take advantage of the intended signal’s physical proximity to the sensor and the ability to elicit feedback to discriminate between real and forged signals.

II. SENSOR AND EMI ATTACK OVERVIEW

A. Threat model

This work considers an adversary that has prior knowledge of the device under attack, including the specific make and model of the device — information that could be




Figure 2. Safe zone (Zone 2) for operation of CIEDs from ANSI/AAMI PC69-2007 [17]. Emissions in Zone 1, even at low amplitudes, have a higher risk of interference because it is the sensing region for those devices.

obtained via other channels including social engineering. For devices in widespread use, the attacker may even possess a device of the same model.

In addition, we assume that an adversary has access to commodity hardware (e.g., laptops, audio amplifiers, and signal generators) sufficient to mount attacks from a distance of several meters. Although the range can be increased with specialized equipment, this work aims to demonstrate the attack feasibility and focuses on techniques that can bypass filters and common defenses; we do not directly address the transmission power. A well-funded adversary could launch longer-range attacks using high power amplifiers and high-gain antennas.

The adversary’s goal is to manipulate sensor readings by injecting signals directly into the analog circuit without altering the sensed physical phenomenon. Thus, if $s(t)$ represents the readings produced by the sensor in isolation, the adversary’s goal is to inject a malicious signal, $m(t)$, such that the sensor readings become $s'(t) = m(t) + s(t)$, where $m(t) \gg s(t)$, and $m(t)$ dominates with $s'(t) \approx m(t)$.

B. Sensor background

Sensors are transducers that convert physical phenomena such as light, temperature, or sound into electrical signals. Cardiac activity produces electrical signals that can be sensed directly. This work considers sensors that produce voltage signals in the Very Low Frequency (VLF) band (1 Hz–30 kHz) or lower. The output is then amplified, possibly filtered, and digitized before delivery to an application running on the microprocessor (Figure 3). The frequency range of the output is the sensor’s *baseband*. Sensors may be sensitive to EMI in their baseband and without filters, they may be sensitive to EMI outside of the baseband too. This work thus divides sensors into two categories: *baseband response* and *high-frequency response*. The rest of this paper examines one type of sensor from each of the above categories.




Figure 3. The typical components of an audio sensing device. EMI attacks target the circuit just before amplification, where the signal is the weakest.

1) *Cardiac medical devices*: Cardiac devices including external electrocardiogram machines (ECGs) and cardiac implantable electrical devices (CIEDs) measure cardiac signals and may deliver therapies as needed. The measured signals, called electrograms, pass through a set of analog filters to remove unwanted frequency components before being amplified and digitized. As a result, medical devices are more resilient against high-frequency interference but may still be sensitive to baseband interference. The ANSI/AAMI PC69-2007 [17] electromagnetic compatibility standard, summarized in Figure 2 indicates that low frequency EM radiation has a significantly greater likelihood to interfere with CIEDs.

The standard separates EMI into three zones. Zone 2 is where many modern devices may emit, including mobile phones [18], [19] and electronic article surveillance (EAS) systems [20]. Zone 1 is the operation and sensing zone for CIEDs which cannot implement aggressive filtering because they would also attenuate the intended signal.

2) *Microphones*: Microphones are part of audio capture circuits and transform acoustic waves into voltage signals. Those signals are then amplified and digitized by an analog-to-digital converter (ADC) before reaching a microprocessor (Figure 3). Audio capture circuits have a baseband ranging from 20 Hz to 20 kHz, but in commodity electronics, they tend to lack filters. In addition, because the expected signal prior to the amplifier is on the order of 1 mV, low-power EMI can cause injected signals to appear in the circuit before the amplification stage. Those signals, if strong enough can dominate the legitimate signals.

C. Manipulating sensor readings

Electromagnetic signals can cause voltage differences to appear across conductors placed in the vicinity. An attacker could use this mechanism to inject unwanted signals into a system. The amplitude of the induced voltage depends on the strength of the electromagnetic field, with low-power EMI typically causing millivolt fluctuations. Digital components are typically well-protected because they operate at multi-volt levels. For example, a microprocessor operating on 2 V can represent the bit '1' with voltages above 1 V, and the bit '0' with voltages below that threshold. Analog sensors, however, are more sensitive to millivolt fluctuations.

To manipulate sensor readings with EMI, an adversary must find a suitable emission frequency. Each circuit component has its own operation frequency band within which




Figure 4. Example EMI on the voltage of an audio signal after amplification on the analog circuit. The electromagnetic interference attack can be high-amplitude baseband (middle) or modulated (bottom). The injected signals dominate after amplification with automatic gain control (right).

a signal pass with little attenuation. Circuit components in series with different operation frequency bands can result in the elimination of a large portion of the frequency bands suitable for signal injection attacks. For example, a short conductor may work well for high-frequency coupling, but a low-pass filter in the downstream path may eliminate all high-frequency components. In such a case, injecting EMI signals in the baseband will likely yield better results.

D. Baseband EMI attacks

Systems that include low-pass filters severely attenuate high frequency signals. Thus, to survive those filters, a malicious injected signal $m(t)$ must be in the baseband: the emitted EMI $v(t)$ must be in the same frequency range as $m(t)$, as shown in the center of Figure 4. Baseband injection requires relatively high power emission because circuits do not normally respond well to radiation in those frequencies.

Some pacemakers can detect cardiac tissue signals by looking for large voltage change rates (slew rates), dv/dt [21] and voltage thresholds [22]. Sending a high-amplitude signal could obscure the actual signal due to automatic threshold or gain control [22]. High-amplitude but benign sources of radiation in those frequencies have been reported to affect some devices [23], [24], but lower-amplitude targeted waveforms have not received much attention.

E. Amplitude-modulated EMI attacks

Amplitude-modulated EMI attacks target systems lacking filters and are thus more likely to respond to high-frequency signals. Circuits may contain components that couple efficiently to signals in the MHz and GHz range. An adversary can thus tune the transmitter to a carrier frequency

(f_c) that closely matches the receiving circuit's resonant frequency and maximize the induced voltage. The baseband injection signal $m(t)$, as a function of time t , can then be upconverted to the carrier using amplitude modulation — much like an AM radio. Thus the modulated EMI signals have the form $v(t) = m(t) \cos(2\pi f_c t)$. On the receiving side, a vulnerable sensor can behave as an AM receiver and downconvert $v(t)$, recovering only the baseband signal $m(t)$. As a result, the output of the sensor is $s'(t) = s(t) + m(t)$. The key to mounting a successful modulated attack is to find a frequency that can induce a large enough voltage $v(t)$, and simultaneously be demodulated by another component to recover $m(t)$ from $v(t)$.

1) *Conducting paths as antennas*: To exploit the “back-door” coupling [1] effect, the frequency of the emitted EMI signal carrier has to be at the resonant frequency of the receiving circuit component in order to maximize the received voltage levels. An approximation to determine the resonant frequency of a whip antenna in far-field communication is its length, which is approximately one quarter of the wavelength of the resonant frequency. This rule may not be applicable to a wire connecting two electric components inside a sensor because the impedance of the connected components is unknown. Thus, the best way to determine the resonant frequency of a sensor and therefore the carrier frequency of EMI signals is to obtain a copy of the device and sweep through a range of frequencies. We call this method *reverse-tuning* and provide details in section IV.

2) *Nonlinear components as demodulators*: In communications, the harmonics and cross-products[†] produced by the nonlinearity of electric components are typically considered undesirable distortions. An adversary can exploit those distortions to achieve downconversion and to obtain the baseband waveform. Ideally the components should be linear devices such as amplifiers that amplify an input signal $v_{in}(t)$ by a gain A . Thus, the output can be described as $v_{out}(t) = Av_{in}(t)$. In practice, amplifiers contain nonlinear components, and the simplest output of a nonlinear amplifier can contain a quadratic term:

$$v_{out}(t) = Av_{in}(t) + Bv_{in}^2(t), \quad (1)$$

where B is the gain for the quadratic term v_{in}^2 . With a crafted input signal, such a nonlinear amplifier can downconvert the signal and recover the baseband signal. For instance, an attacker with the goal of injecting $m(t)$ can induce the following voltage signal as the input to the amplifier,

$$v_{in}(t) = m(t) \cos(2\pi f_c t) + \cos(2\pi f_c t). \quad (2)$$

Without loss of generality, let $m(t)$ be a simple tone, i.e., $m(t) = \cos(2\pi f_m t)$. After applying Eq. (2) to Eq. (1)

[†]Harmonics are frequencies that are integer multiples of the fundamental frequency components, and cross-products are multiplicative or conjunctive combinations of harmonics and fundamental frequency components.

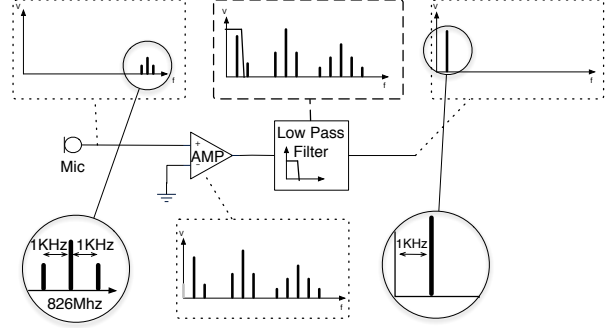


Figure 5. Example where the low-frequency signal $m(t)$ is a 1kHz sine wave modulated on a high frequency ($f_c = 826MHz$) carrier. The injected signal appears on the conductor between the microphone and the amplifier. The nonlinear component (e.g. the amplifier) introduces several frequency components in the baseband and high frequency bands. After the low-pass filter, only $m(t)$ will be left and will be perceived as a real signal.

and taking the Fourier transform, we can confirm that the output of the amplifier contains the intended frequency component f_m together with the amplified fundamental frequency components of v_{in} (i.e., $f_c - f_m$, $f_c + f_m$, and f_c), harmonics, and other cross products (i.e., $f_m, 2(f_c - f_m), 2(f_c + f_m), 2f_c, 2f_c + f_m$, and $2f_c - f_m$), as shown in Figure 5. After a low-pass filter, all high-frequency components will be removed and the f_m frequency component will remain, which completes the downconversion.

3) *Analog to digital converters as demodulators*: During the digitization process, an ADC with a given sampling frequency is used. By matching the emitted EMI carrier frequency to the sampling frequency of the ADC an attacker can turn it into a demodulator. Specifically, to yield a digitized and discrete sequence $v[k]$, an ADC samples a continuous analog signal $v(t)$ every T_s seconds, i.e., $v[k] = v(t)|_{t=kT_s}, k \in [1..\infty]$. Let $V(f)$ be the Fourier transform of the original analog signal, and let $V_s(f)$ be the sampled signal. Then,

$$V_s(f) = f_s \sum_{n=-\infty}^{\infty} V(f - nf_s),$$

where f_s is the sampling frequency. Essentially, the sampling process creates a duplicated spectrum of the original signal by shifting to $f - nf_s$ for $n = [-\infty..1, 0, .. + \infty]$. An adversary can select the carrier f_c to be a multiple of the sampling frequency f_s , e.g., $f_c = 9f_s$. Thus, during the digitization process, the ADC will sample the carrier at intervals that skip the high-frequency oscillation, thus acting as an envelope detector and recovering the original $m(t)$.

4) *Capacitor and diode as demodulators*: The envelope of the attacking signal in Eq. (2) is given by

$$e(t) = |m(t) + 1|.$$

If $m(t) + 1 \geq 0, \forall t$, then the modulated EMI signal $v(t)$ can be demodulated by passing through a simple envelope

detector consisting of a diode and a capacitor, which also happen to be a basic building blocks in many circuits. In at least one of the devices we used (the MTS300CB board for MicaZ motes) we discovered that there are several capacitor combinations on the path between the microphone and the amplifier. While a diode was not present in that particular circuit, the amplifier forces the current to flow in one direction through the circuit, thus having the circuit itself behave as a diode and making that circuit a good candidate for extracting the baseband signal $m(t)$.

F. Distance bounds

An important factor is the relationship between the feasible attack distance and the strength of an injected electric signal. In a receiving circuit with resistance R_r , in order to induce an electric signal with V volts, the received power (denoted by P_r) of an EMI signal is

$$P_r = \frac{V^2}{R_r}. \quad (3)$$

Assuming that modulated EMI signals typically operate in the MHz–GHz frequency band and that adversaries are at least 0.5 meters away, we formulate the signal propagation as far-field communication. Although an accurate radio propagation model should account for multipath, shadowing, and fading, we utilize the free space propagation model to understand the basics of feasible attack distances. Consider an attacker at a distance d from the victim’s circuit who transmits at a power level of P_t . Then the received power, P_r , is calculated from the Friis transmission equation.

$$P_r = P_t G_t G_r \left(\frac{\lambda}{4\pi d}\right)^2, \quad (4)$$

where G_t and G_r are the antenna gains of the transmitting and receiving antennas respectively, and λ is the wavelength of the signal.

For example, consider an adversary who transmits at 100 mW with a 10 dB antenna. Suppose that the victim’s device responds well to an 826 MHz carrier with a receiving circuit of resistance 1.5 k Ω and an extremely low gain of 0.01 dB since it was not designed to receive radio signals. To induce 10 mV on the victim’s sensing circuit during an attack, from eq. (4) and eq. (3) the distance between the attacker’s antenna and the victim’s system can be at most 11.2 m; this makes the attack practical. Using an antenna with a higher gain of 20 dB and a signal source output of 1 W could increase the attack distance to over 50 m.

III. BASEBAND ATTACK METHODS AND EXPERIMENTS

Safety-critical systems such as medical devices commonly have low-pass filters that attenuate high-frequency signals where the resonant frequency ranges reside. An attacker can still send low frequency signals within the passband of the filters and compensate for the frequency mismatch with higher-power signals (around 1 W or more) or reduced distance.




Figure 6. Generated forged heart beat with recognizable peaks at 1.1 Hz (Top). Print-out of the electrocardiogram of the patient simulator configured to exhibit ventricular fibrillation (Bottom). The induced signal is visible on lead 5 in the middle right. The ECG erroneously reported a 66 bpm pulse.

A. External ECG

An electrocardiogram (ECG) device is designed to monitor cardiac activity (around 1 mV) by taking voltage readings at the skin surface. ECGs connect to patients with conductive pads and leads.

1) *Experimental setup:* To investigate the effects of EMI on the system, we plugged an ECG to a patient simulator (Bio-Tek Lionheart) configured to exhibit symptoms of ventricular fibrillation. We used an arbitrary function generator (AFG) connected with a simple whip monopole antenna to radiate low-power EMI signals. To compensate for the low transmission power and the inefficient radiator, we left two of the leads disconnected and placed them within 5 cm of the radiating antenna. Finally, we read the ECG screens to determine whether the EMI injection was successful.

2) *Antenna and vulnerable frequency range:* When sending high-frequency signals above 1 MHz, no signals were observed at the ECG, although an induced voltage of over 10 mV_{p-p} was measured at the leads. The results showed that the pads and leads can serve as the entry point for EMI to alter the sensor readings, but the high frequency attenuation would make modulated EMI attacks difficult.

3) *Baseband EMI attacks:* We transmitted a baseband signal that emulated a cardiac rhythm (heart beats) at 66 bpm, as shown in Figure 6 (Top). After a stabilization period of 60–120 seconds, we observed that on the right side of the printout sheet in Figure 6, the peaks of the injected baseband signals are visible, indicating that our EMI signal affected the sensor readings.

B. Cardiac implantable electrical devices (CIEDs)

CIEDs are used to treat cardiac diseases with electrical stimulation. Under most configurations, pacemakers and defibrillators will send low-energy electrical stimulations (around 10 μ J) to pace the cardiac tissue if no cardiac activity



Figure 7. Illustration of the connector and tip of typical active fixation corkscrew bipolar pacemaker leads. The central cathode connector is extended further than the external anode and the cathode tip protrudes further than the anode ring.

is detected. Additionally defibrillators can be configured to detect potentially dangerous rhythms such as fibrillation or tachycardia, and deliver a shock (around 25 J) [25] to reset all the electrical cardiac signals so that normal cardiac activity may resume.

1) *EMI coupling on leads:* Leads threaded through blood vessels and into the cardiac chambers connect a CIED to the cardiac tissue. With the standard lead design, it is possible to induce a voltage when exposed to EMI [20]. Due to flexibility requirements, instead of true coaxial designs, leads use wound coaxial designs with the outside conductor (i.e. the anode) tightly wound around the central insulated conductor (i.e. the cathode). On the lead end that connects to cardiac tissues, the central wire protrudes at what is called the cathode tip. That tip is about 2.4 cm longer than the external conductor that stops at what is called the anode ring (Figure 7). On the lead end that connects to the CIED, the anode is 1 cm longer than the cathode. This design combined with the difference in length between the anode and cathode conductors allow a voltage to be induced by EMI. When radiating a 32 MHz signal at 300 mW from 14 m away, we observed induced voltages of around 100 mV between the anode and the cathode of a bipolar lead. Signals under 1 MHz induced lower voltages.

2) *Vulnerable frequency range:* CIEDs are designed to amplify specific regions such as the 0.1 Hz–1 kHz range [23] that contain electrogram information. To study their frequency response, we first disassembled an implantable pacemaker and measured the output of the first filter that is connected to the lead. The measurements show that signals below 5 MHz are only attenuated by 4 dB, and signals beyond 200 MHz and 800 MHz are attenuated by 30 dB and 40 dB respectively, making high-frequency signals a poor choice for an attacker. To obtain the system wide frequency response, we ran a sweep in the low-frequency range between 0.1 Hz and 1 kHz and observed the resulting relative signal amplitude reported on the programmer connected to the device. Signals in the 100 Hz–300 Hz range from our amplifier showed the strongest amplitudes.

3) *Baseband attack experimental setup:* We performed the experiments under different conditions including:

- Free air, providing conditions to find good candidate waveforms for injected signals that would have a measurable impact on the system;

- Saline bath with a 1.8g/L NaCl concentration, built following the ANSI/AAMI PC69 specifications for electromagnetic compatibility testing of cardiac devices [17]; and
- Synthetic human with a functioning circulatory system (using saline solution) and partial model of the human heart [26]. Resistance measurements of the synthetic human’s tissue showed that it is an approximate match with human tissue.

In our experiments we used 3 defibrillators (Medtronic InSync Sentry - 2005, Boston Scientific Cognis 100-D - unknown year, St. Jude Promote - 2007) and one pacemaker (Medtronic Adapta - 2006). For each CIED tested, we attached the same set of bipolar sensing and pacing leads (Pacetrionix Model No. 3851 VB) in the Left Ventricle (“LV”) port. Preliminary results with a different set of leads (Guidant Dextrus 4137) showed comparable results. For the attacking waveform frequency in the 0.1 Hz to 1 kHz range, we used an audio amplifier connected to a wire used as a simple whip antenna. The amplifier’s estimated output was 10 W, corresponding to 50 volts over a 250 Ω load at the output. The effective radiated power was much lower because of the mismatched antenna. Nevertheless, our radiating system was sufficient to produce noticeable induced signals.

Our goal was to create pacing inhibition (atrial or ventricular) and defibrillation shocks. To determine success, we used the electrogram readings displayed by a programmer compatible with the device under test. To inhibit pacing, we injected a 100 Hz sinusoid signal, a pulsed sinusoid with a 100 ms width at 1 pulse per second, and a waveform from the ANSI/AAMI EC13 set amplitude-modulated over 100 Hz. To induce fibrillation events, we used the waveform number 421 from the MIT-BIH Malignant Ventricular Ectopy Database, an electrogram recording of a real episode of ventricular fibrillation.

4) *Results:* The results are summarized in Table I. Pacing inhibition in free air can be accomplished from the furthest distance (1.5 m), and pacing inhibition in saline solution were the most difficult condition to produce (5 cm or under). In free air, we did 20 trials for all devices except for the St. Jude device where only 8 trials were recorded. The synthetic human measurements only had 2 trials.

Pacing inhibition. In the free air tests, the pulsed and modulated sinusoid signals at 30 V p–p effectively stopped pacing on all tested devices from 0.68 m to 1.57 m. A sample output from the Medtronic programmer is shown in Figure 8 (Top) showing a purple Ventricular Sense (VS) marker after the onset of our EMI signal.

Testing using saline solution, we used two setups. We first arranged the leads in an arc and completely submerged the device in the saline solution. We injected signals with a maximum amplitude of 50 V p–p, but we were unable to cause pacing inhibition with submerged devices, although our attenuated signal was observed on the EGM. In a second

Device	Open air	Open air (defibrillation)	Saline bath	Saline (lead tips only)	SynDaver
Medtronic Adapta	1.40 m	<i>Not applicable</i>	<i>No inhibition</i>	0.03 m	<i>Untested</i>
Medtronic InSync Sentry	1.57 m	1.67 m	<i>No inhibition</i>	0.05 m	0.08 m
Boston Scientific Cognis 100-D	1.34 m	<i>No defibrillation</i>	<i>No inhibition</i>	<i>Untested</i>	<i>Untested</i>
St. Jude Promote	0.68 m	<i>No defibrillation</i>	<i>No inhibition</i>	<i>Untested</i>	<i>Untested</i>

Table I

MEDIAN MAXIMUM DISTANCE AT WHICH A REPOSE (PACING INHIBITION UNLESS OTHERWISE SPECIFIED) FROM THE DEVICE WAS OBSERVED. DEVICES WERE CONFIGURED IN BIPOLAR MODE. DIFFERENT WAVEFORMS WERE USED FOR THE PACING INHIBITION AND DEFIBRILLATION TESTS.

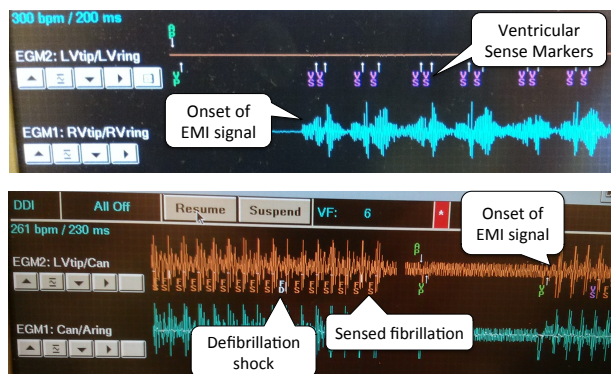


Figure 8. Free air interference with the Medtronic InSync Sentry implantable cardiac defibrillator. After the onset of the EMI signal, the device reported ventricular sense (Top) and fibrillation sense (Bottom). The FD label indicates the point at which the defibrillation shock was delivered.

setup with only the bipolar leads tips in the solution and the device in free air, configured to sense in bipolar mode, we can inhibit pacing at a range of 2 cm to 3 cm.

We also tested the the Medtronic InSync Sentry implanted in the synthetic human. For these tests, a cardiologist used the common approach by threading the leads through the axillary vein under the left clavicle [27]. The tip of the lead was guided inside the model heart. With the saline solution flowing through the model at approximately 60 bpm, we were able to inhibit pacing only using the modulated sinusoid signal with an amplitude of 50 V p-p at a range not exceeding 8 cm from the leads.

Defibrillation shocks. For the fibrillation test, we used the Medtronic InSync Sentry in free air and saline, along with the 421 waveform as mentioned above. A snapshot of the response as displayed by the programmer is shown in Figure 8 (Bottom) with markers showing fibrillation sense (FS) events and a defibrillation (FD) event after the onset of the EMI signal. With our 10 W amplifier, the median maximum range was 1.67 m in free air, and the results were negative in saline solution. These results indicate that an attack is possible with the waveform we used, but a power source greater than 10 W would be needed.

IV. MODULATED ATTACK METHODS AND EXPERIMENTS

Devices lacking low-pass filters are more sensitive to EMI because they do not attenuate signals outside of the baseband, and thus can act as efficient receivers for high-frequency signals close to the circuit’s resonant frequency. This work focuses on microphones as example of unfiltered devices; first outlining a reverse AM tuning method to locate the resonant frequency and then demonstrating the injection of: DTMF tones, music, speech, and audio test waveforms. We use the Speech Transmission Index [28] and the Shazam service based on spectral fingerprinting [29] to evaluate the strength and fidelity of injected signals.

A. Finding the resonant frequency

The conducting path between a microphone and the accompanying amplification circuit can act as an antenna, as discussed in Section II. This creates a likely entry point for signal injection. To launch a successful attack, an adversary must find a frequency satisfying two conditions on the target circuit: (1) suitable for demodulation of the baseband signal, (2) close enough to the resonant frequency to induce a high voltage. It is difficult to calculate the frequency responses of the conducting path and other circuit components, especially if no technical details are available. However, an attacker can measure the resonant frequency empirically, possibly with partial information of the device, such as the length of the conducting path.

Based on the technical details available and analysis necessary to develop an injected waveform, modulated EMI attacks fall into three categories: black box, gray box, and white box. We tested each class of attack using a signal generator that operates in the 9.00 kHz–2.02 GHz frequency range to modulate and transmit signals. The baseband waveforms used include a simple 440 Hz sinusoid, and an arbitrary audio waveform called the “Weezer” waveform after the band that produced the sample [30].

1) *Black box with no technical details:* We first used a webcam (Logitech Quickcam Ultravision) with a camera and microphone integrated into a single enclosure. With no directly observable indication of the conducting path length between the microphone and the amplifier, we could not approximate the resonant frequency. For all experiments, we connected the webcam to a laptop to capture the audio output




Figure 9. Power spectral density of the recovered 440Hz tone modulated on a range of carrier frequencies. The 820 MHz to 840 MHz range showed a high response, indicating that carriers within that band are likely to successfully inject signals.

and fixed the signal generator’s output power at 80 mW. For short-range experiments, we used a whip antenna placed within 20 cm of the webcam. For longer-range experiments, we used a dipole antenna with higher gain to obtain the same results from ~ 1 m away.

Results: Figure 9 shows the amplitude of the received signal when modulating a single 440 Hz tone onto a range of carrier frequencies. This test revealed that the webcam’s resonant frequency was between 820 MHz and 840 MHz. We next attempted to inject an arbitrary waveform by modulating the Weezer waveform onto a frequency in the resonant range. The commercial Shazam service, which uses a spectral fingerprinting method to identify audio samples, correctly identified the recovered waveform, indicating that the audio was clear enough to be recognized.

2) *Gray box with limited technical details:* We next found the resonant frequency of a Bluetooth headset (Plantronics Voyager 510) based on partial information of the circuit. A disassembly revealed that the microphone–amplifier conducting path was approximately 6 cm, which corresponds to a resonant frequency of approximately 1.25 GHz if we model the conducting path as a whip antenna. Thus, instead of the entire frequency range, we swept the carrier frequency from 1.0 to 1.5 GHz to pinpoint the resonant frequency. The output of the signal generator was fixed at 20 mW and we positioned the transmitting antenna 10 cm from the headset.

Results: Using the single tone at 440 Hz, we found the resonant frequency of the headset at 1.175 GHz, which matches the predicted range. In this case, the length of the conducting path was useful in locating the resonant frequency. We then modulated the Weezer waveform onto the resonant frequency and the recovered signal was correctly identified by the Shazam service.

3) *White box with available technical details:* A MicaZ mote fitted with an MTS300CB sensor board served as a white box attack target. The manufacturer’s schematic documents multiple capacitors in the path between the microphone and the amplifier. These components are analogous to a simple envelope detector, as discussed in Section II-E, and are a likely entry point for injected EMI. To ensure that the resonant frequency was within the limits of our

transmitting equipment, we modified the circuit by fitting 15 cm wires between the microphone and the rest of the board. We then measured the induced voltage on the leads as we swept the EMI signals from 9.00 kHz to 2.02 GHz with the output at 20 mW and the transmitting antenna less than 1 m from the the MicaZ mote. The mote digitized the measured audio from the sensor board and forwarded the data to a laptop via a second mote acting as a base station.

Results: By first measuring the voltage at the wires, we noted a peak around 83 MHz and confirmed the suitability of this carrier frequency for signal injection by modulating a single 440 Hz tone and recovering it at the application layer. When we tested the Weezer waveform the recovered waveform was recognizable, but the Shazam service was unable to retrieve the record, possibly due to the change in pitch and speed resulting from the mote’s primitive codec. We also noticed that an efficient whip antenna for 83 MHz should be about 6 times longer than the wires we used. This mismatch suggests that estimating the resonant frequency based only on the length of the conducting path may not be a reliable technique. The unknown impedance of the receiving circuit may alter the expected resonant frequency sufficiently to require a manual frequency sweep in some cases.

B. Dominating a legitimate signal

It is difficult to inject forged signals that can remove legitimate signals because cancellation requires a high-fidelity model for the waveforms arriving at the sensor. Instead, an attacker can inject a powerful forged signal to dominate the legitimate signal. In systems with automatic gain control, powerful injected signals could force the gain to be automatically reduced to avoid circuit saturation. As a result, the legitimate signal experiences fading.

There are two possible outcomes for an attacker attempting to overwhelm a legitimate signal: (1) The legitimate signal is low and the injected signal dominates, leading to a successful attack; (2) The legitimate signal is high and the injected signal cannot completely dominate without saturating the amplifier. In this case, the attack acts as a simple denial of service resulting in distorted audio.

To quantify the effectiveness of modulated EMI attacks, we use the Signal-to-Interference Ratio (SIR) defined as

$$SIR = 10 \times \log_{10} \left(\frac{P_{Signal}}{P_{Interference}} \right) \quad (5)$$

where P_{Signal} and $P_{Interference}$ are the power levels of the measured signal and the induced forged signal respectively. The SIR quantifies how much stronger the legitimate signal is relative to the injected signal. A negative SIR indicates that the injected signal is stronger than the legitimate one and the legitimate signal is difficult to recognize.

Results. We used an *audio* tone (440 Hz) as the legitimate signal and a single tone (550 Hz) modulated over the resonant 826 MHz carrier as the EMI signal. Figure 10 shows the




Figure 10. Signal-to-Interference Ratio (SIR) as the transmission power is increased with $50\ \Omega$ at the output. At an SIR below 0 dB, the injected signal dominates. The 3 sharp dips were caused by the transmitting signal generator switching between power modules.

SIR when varying the output power of the signal generator. As expected, the SIR has an inverse relationship with the output power of the signal generator. In the region where the $SIR < 0$, the interfering signal dominates.

C. Transmitting intelligible speech via EMI

To determine the feasibility of reliable intelligible speech transmission over EMI, we used the *Speech Transmission Index* (STI) [28], a standard measure to predict speech intelligibility. The index is computed from the signal-to-noise ratio (SNR) at 8 octave bands covering the range from 125 Hz to 8 kHz using the formula $STI = \sum_{i=1}^n W_i [\frac{SNR_i + 15}{30}]$, where $n = 8$ for most applications, W_i is a predefined weight assigned to each octave band, and SNR_i is the received audio SNR at the octave band i . The STI ranges from 0 to 1, where 1 indicates a high likelihood of intelligible speech transmission. We compared results between the injected audio over EMI and the legitimate acoustic channel.

Experimental setup. We chose three devices with microphones: the Bluetooth headset and webcam used earlier, in addition to another webcam, a Logitech Quickcam Vision Pro. For the audio, we used a standard STI test waveform consisting of a uniform mixture of frequencies to allow an objective assessment of the response profiles. For the *acoustic channel*, the STI waveform was played through a MacBook Pro speaker system and recorded by the microphone of the device under test. For the *EMI channel*, we modulated the STI waveform on the resonant frequencies of each device and transmitted the result over the air.

Results. We computed the average STI for the 3 devices over audio with a mean of 0.69, and over EMI with a mean of 0.72, indicating that both channels are comparable for speech transmission. Notice that the waveform transmitted via the EMI channel has a slightly higher STI. The injected signal induces voltages on the conducting path between the microphone and the amplifier and is therefore free of the mechanical limits of the microphone itself. That effect is also apparent in the lower octave bands in Figure 11.

D. Ghost talk use cases

We tested the modulated EMI attacks in a few real-world scenarios: an automated telephone system, audio phone calls, and video teleconference calls. For all scenarios, we tested three cases: a Bluetooth headset paired with a phone that made calls over cellular networks, a Bluetooth headset paired with a laptop that made calls over VoIP (using Skype or Google Chat), and a webcam connected to a laptop that made calls over VoIP.

Automated dial-in system. Many automated dial-in systems take their customers' inputs via telephone DTMF (Dual Tone Multiple Frequency) signals — sending keypad presses as a unique combination of two audio tones. To demonstrate reliable transmission of DTMF tones via EMI, we connected to the dial-in service of Citibank's credit card system. We then successfully entered the credit card number and zip code sequences by injecting the corresponding DTMF tones via modulated EMI attacks, giving us access to the credit card information. This result shows that it is possible to use EMI signals to initiate virtual button presses via DTMF during a victim's phone session with a remote system.

DoS attacks. To determine if we could overwhelm acoustic signals to the point where none of the original signal was apparent, we used the Weezer waveform and increased the power as far as possible without causing distortion in the demodulated audio. We then mounted an attack against a Skype session initiated with the Bluetooth headset. Shortly after the conversation started, we began transmitting the EMI signal. The injected signals overwhelmed the acoustic signal to the point where the remote user could not detect the original acoustic signal. This result demonstrates that blocking a legitimate conversation is possible. For better results, a sound-masking noise, such as a source of white noise with a uniform continuous spectrum could be used [31]. We had similar results with a webcam connected to a laptop, as well as transmission at higher power.

Session hijacking. Instead of completely blocking the acoustic signal, we also tried substituting the acoustic speech signal with an EMI speech signal. This attack is similar in principle to injecting music, but the receiving user instead hears speech that could plausibly replace the caller. We transmitted an EMI signal modulated with a reading of *Edgar Allan Poe: The Raven* by James Earl Jones. We transmitted the signal immediately after the victim initiated a phone call using the Bluetooth headset. The injected speech introduced additive audio signals observable by the calling party, but it did not completely mask the victim's voice. To the receiving user, the acoustic signal appeared as background noise with the EMI signal coming across clearly. We suspect that it is possible to obscure the victim's voice as long as the EMI signal is powerful enough. In our case, we were limited by the signal generator with a maximum output power of 80 mW and a low-gain antenna.




Figure 11. Frequency responses of acoustic input and EMI signal for the devices tested using the Speech Transmission Index waveform. The high SNR over EMI at low frequencies indicates a better response than the audio signal, possibly due to bypassing mechanical constraints in the microphone.

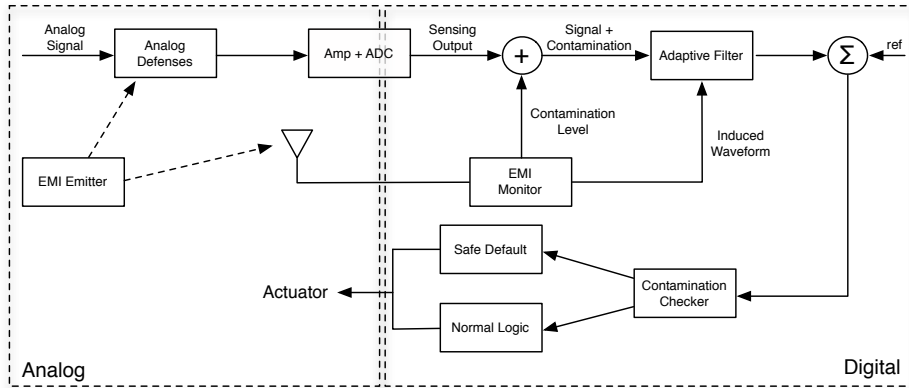


Figure 12. Overview of the defenses, integrated in a single system. The signal first goes through a number of analog defenses to attenuate the induced signal before conversion into a digital format. A subsystem simultaneously takes readings of the EMI level in the environment to determine the contamination level. Further filtering is then possible based on the estimated EMI level. Active probing can also be used to help discriminate between induced and measured signals. Finally, if the probing results indicate a forged waveform capable of forcing improper actuation, we revert to a safe default.

V. DEFENSES

A. Analog defenses

The goal of our defense set is to improve the trustworthiness of the sensor readings by attenuating the induced signals or at minimum detect EMI attacks. We propose a system (Figure 12) composed of a series of analog and digital defenses that can attenuate the EMI on the analog sensor circuit, differentiate between induced and measured signals in the digital circuit, remove the induced signals if possible, and revert to known safe defaults if the interfering signal is too strong. Although the analog defenses are known and some are already applied to implantable medical devices, consumer electronics are less protected. As the cost of deploying those defenses is device-specific, we instead quantified the attenuation which can be used for the cost/benefit analysis. The analog defenses on their own may not be enough against strong emitters or baseband emitters. Thus in addition, we propose to use some digital defenses: adaptive finite impulse response filters that can improve the SIR, a probing-based method that can distinguish between induced and measured signals, and a safe default mode for devices in the presence of strong EMI attackers.

In the analog portion of the sensing circuit, there are three common defenses: shielding, differential comparators, and filters. All of these are used to some extent in modern CIEDs, but not in commodity electronics. In this section we apply those techniques and measure the resulting attenuation with waveforms used in our modulated attacks. Those results can then be used for cost/benefit analysis to improve the design of current systems.

1) *Shielding:* The application of a conducting material to shield a component from electromagnetic radiation is well-known but absent from most commodity devices we tested. We coated the exterior of one of our webcams with a conducting surface leaving large holes for a number of components. Those include the camera lens, two large buttons on the side, the microphone and the mechanical stand. Even with large imperfections in the shield, the attenuation of the recovered EMI signal was over 40 dB, forcing an attacker to transmit 10^4 times more power to have the same effect.

2) *Differential comparator*: Where shielding is either not possible or not sufficient, a reference signal can be used to remove the common mode voltage using a differential circuit, commonly used in analog electronics [32]. By measuring the difference in potential between two voltages, the common mode interference present on both signals is effectively filtered out.

Early designs of pacemakers used unipolar leads (one single conductor in lead connected to the cardiac tissue), which was eventually phased out in favor of the “true bipolar” design [27]. Under similar conditions with a signal injection waveform (see Section III) we measured the attenuation of the induced signal from a bipolar lead to be 30 dB, showing a significant reduction in EMI induced signal with a simple defense.

Even with a bipolar design, our results indicate a possibility to induce a differential mode signal across the anode and the cathode of the leads. In free air, with a 100 Hz sinusoid on an 80 mW source at 20 cm, we measured that the induced voltage difference is on the order of 5 mV and the phase shift is on the order of $\pi/10$. The differential voltage and phase difference became very small with the tips dipped in the saline solution, leading to a severe drop in the differential voltage. We had to significantly increase the transmit power of our emitter to return the measured differential voltages to around 5 mV even when dipped in saline solution, indicating that the attacker can compensate with increased power.

3) *Filters*: A filter that attenuates signals outside a sensor’s baseband frequency can reduce the vulnerable frequency range of that sensor. Such filters are already in use in medical devices, but they seem more sparse in commodity electronics. Those are therefore more vulnerable to signal injection attacks with a high frequency carrier, better matched to the vulnerable circuit. To test the effectiveness of filters for commodity electronics, we used a custom-built active low-pass filter at 500MHz to attenuate high-frequency components while allowing audio signals (below 50 kHz) to pass. In the case of the 836 MHz carrier from the attack waveform we suggested in Section IV, an attenuation of over 40dB was measured, making it a very good attenuator against our signal.

B. Digital Defenses

Due to physical requirements or packaging limitations of implantable medical devices, some of the analog defenses outlined above may have a limited effect, especially against a strong emitter. As a result, the output of the sensor may still contain injected signals. In addition to the analog attenuation defenses, we propose techniques on the digitized signal to estimate, track, clean, and verify the state of the signal as it moves through the system.

1) *Signal contamination*: A necessary component of our signal injection attack is an EM wave to carry the signal. If a component in the victim’s device is available to capture

only the radiated signal, we can estimate the EMI level in the environment. We call this estimation the *signal contamination*. Components downstream can then use the contamination level to determine the appropriate defenses to apply.

As a metric of the required conditions for EMI attacks, we use the root mean square of the waveform amplitude in a window of size w ms to estimate the EMI level in the environment. To compute the contamination level, l_c , we compare the measured level (A_t) to one calibrated in a quiet environment (A_0).

$$l_c = \frac{RMS(A_t)}{RMS(A_0)}, \quad (6)$$

where t is the start time for a window of size w under consideration, and RMS is the root mean square of the waveform amplitude as defined by

$$RMS(A_T) = \sqrt{\frac{1}{w} \int_T^{T+w} A_t^2}, \quad (7)$$

where $T < t < T + w$.

In communicating implantable medical devices, the RF antenna can be used to estimate the ambient EMI. In non-communicating devices that lack an antenna, a reference conductor can be used instead. We assume that the monitoring component is located close enough to the vulnerable part of the sensor to receive comparable levels of EMI radiation. In the presence of a pulsed EMI signal with increasing power, similar to the pulse from Section III, the communication antenna on our disassembled CIED recovered the signal shown in the spectrogram in Figure 13 (Top). The 300Hz pulses with increasing power are clearly visible. The bottom graph shows the computed contamination level with a window of size $w = 100$ ms based on the observed waveforms. In our free air experiments, a pacing inhibition would occur for devices exposed to contamination levels of 2.1 and higher.




Figure 13. Top: Increasing strength of attacker signal. Bottom: Computed contamination level.

2) *Adaptive filtering*: If the contamination level exceeds a threshold that can cause improper actuation, we can activate an adaptive filtering mechanism [33] and use the measured contamination waveform to estimate the RF-induced voltage at the leads and clean the received signal. The adaptive filters dynamically adjust the signals between the leads and the antenna to determine a map that can be used to translate the waveform between the two components[†]. Our end goal is to attenuate the EMI-induced waveform on the signal at the leads, thus increasing the Signal-to-Interference ratio.




Figure 14. Adaptive noise canceling system adapted to cancel the forged signal. The output of the system is used to estimate the forged signal given the waveform on the compensating circuit.

Figure 14 summarizes a typical design for an adaptive noise cancelling system [33]. The sensing circuit and the compensating circuit act differently to the forged signal $m[n]$ added to them. $y[n]$ is the observed signal composed of the measured signal $x[n]$ added to some noise $w_1[n]$ and the forged signal $m[n]$. We assume that the antenna component is a linear system of otherwise unknown properties. The output from that circuit is $v_2[n]$ and feeds into the adaptive Finite Impulse Response (FIR) filter. That filter uses the output from the adaptive algorithm to estimate the original forged waveform $\hat{m}[n]$. The resulting waveform is subtracted from the observed waveform at the lead to yield the cleaned-up output which feeds back into the adaptive algorithm to allow the system to adapt to changing waveform amplitude. The reaction time of the filter depends on coefficients used in the adaptive algorithm.

We tested the algorithm against one of our most effective attacking waveforms from Section III, namely the EC13 modulated on a 100 Hz sinusoid. The results shown in Figure 15 indicate a large error at the onset of the attacker's signal (top plot), but the error quickly decreases allowing a recovery of the original measured waveform (bottom plot). The forged signal in this case was severely attenuated, leaving a relatively clean measured signal that would be otherwise obscured.

3) *Cardiac probe*: Systems that can measure the result of their actuation may be able to distinguish if they are under attack. The basic idea is to use the actuation to determine if the sensor readings follow the expected readings. If the attacker cannot observe the actuation, the advantage is further tipped towards the victim.

[†]The mapping accounts for differences in the locations, impedance, and shapes of the two antennas




Figure 15. FIR filters against one of our most effective attacking waveforms. The measured signal waveform is the original ANSI/AAMI EC13 fig 3a, used to simulate a human heart beat. We observe that the estimation error is a maximum in the first cycles of a large amplitude change, but the system quickly adapts, reducing the level of the forged signal, thereby increasing the SIR and the probability that the system would behave correctly in the presence of this waveform.

For CIEDs, in the presence of a sufficiently powerful attacker, the signal received at the processing point in a cardiac pacing device may still contain a residual interference with a contamination level above 2.1 that could cause improper actuation. To discern between the measured and induced waveforms, a CIED can use its direct connection to the cardiac tissue to test if the signal was legitimate.

During a normal cardiac cycle, the tissue contraction produces voltage peaks observable on electrograms measured by the intra-cardiac implanted leads. Immediately after a contraction, during a brief time span called the *absolute refractory period* (ARP), the cardiac tissue will not contract again, even if stimulated with a low-energy pacing pulse (around $10 \mu\text{J}$.) We use this property to discern between real and forged signal. We send a pacing pulse immediately after detecting a voltage peak that may result from a contraction in the tissue around the lead tip. That pulse should reach the cardiac tissue while it is still in the ARP and therefore we expect no signal back on the lead for about 200 ms [22]. If we observe another peak immediately after sending our pacing pulse, there are two possible causes:

1. The cardiac tissue contracted, indicating that a forged peak was present in the signals; or
2. Independent of the tissue response, there was an induced peak on the lead.

In either case, the signal from the lead is not trustworthy and the sensing signal should not be used.

To test the response of the cardiac tissue to our pacing pulse probe method, we used the University of Pennsylvania's Virtual Heart Model [34], [35]. We set the simulator to the default cardiac configuration with no running pacemaker and the intra-cardiac monitoring probe and pacing lead at the tip of the right ventricle. After obtaining a stable cardiac rhythm, we sent pulses both during the absolute

refractory period (ARP) (Figure 16 Middle) and after the ARP (Figure 16 Bottom). The 10 mV response from the cardiac tissue pulse was observed within 40 ms of the onset of the 10 ms pacing pulse. These results suggest that a pacing pulse could give us information on the current state of the cardiac tissue. However, more studies are necessary on the health care aspect of this proposed method.




Figure 16. Cardiac tissue response to pacing pulses. Top: No pacing. Middle: Pacing pulse during the absolute refractory period. Bottom: Pacing pulse after the absolute refractory period. Note that the cardiac wave happens right after our pacing pulse.

During a cardiac probe test, it is possible for an attacker to attempt to cheat the test by sending a high-amplitude pulse during the absolute refractory period, even though the real cardiac tissue is not reacting. Since the attack is blind, the attacker doesn't see the actual cardiac signals; consequently aiming for the ARP will essentially be a random hit. Therefore the probability of landing in the ARP is $Pr[ARP] = \frac{t_{arp}}{t_{interval}}$, where t_{arp} is the absolute refractory period, and $t_{interval}$ is the beat-to-beat interval. To reduce the attacker's probability of success, it is possible to do multiple probes, forcing an attacker to be repeatedly successful in sending pulses in the ARP to avoid detection. To inhibit pacing successfully, the forged pulses need to be sent continuously. The probability of success thus decreases with the number of probes, but not as fast as having the pulses be independent events. Thus, the probability of avoiding detection has a lower bound of $Pr[ARP]^n$, where n is the number of probes used.

4) *Reverting to a Safe Default*: If the previous test determines that the signal is not trustworthy, the system has three options: 1. Disconnect the output from the input; 2. Limit the output to a known safe range; or 3. remove the victim from the environment. In the first case, we can revert to asynchronous pacing – Atrial (AOO), Ventricular (VOO) or Dual (DOO) – as programmed by the medical personnel. In the second case, the output can depend on the

sensing input, but only limited to a safe predetermined range. That technique is heavily dependent on the sensor and the application and is left as future work on specific systems. In the third case, the victim can be notified through an audible alarm about the possibility of an attack, and allow the victim to be moved away from the attacking emitter.

VI. RELATED WORK

The effects of electromagnetic interference either received on electronic circuits or emitted from them are known, although the intentional injection of forged signals due to “back-door” coupling on the analog sensing circuit remains to be investigated. Anderson [6] outlines situations allowing signal leakage from electronic devices, and describes methods for information exfiltration due to EMI sometimes referred to as *TEMPEST* in the military milieu. Further, the U.S. National Institute for Technology and Standards (NIST) issued the FIPS 140-2 document [36] that specifies defense techniques against electromagnetic radiation exfiltration for unintentional radiators and digital devices conformant to requirements from 47 Code of Federal Regulations, Part 15, Subpart B, Class A and B. We didn't expect commodity electronics to be compliant with FIPS 140-2; however we were surprised at how easy it was to penetrate the device's analog sensing system with rogue signals. Given how commodity and COTS electronics are on today's critical path, it is essential to understand the extent of the vulnerabilities on the analog sensing circuit.

Electromagnetic compatibility. The field of electromagnetic compatibility (EMC) is largely devoted to avoiding, minimizing or coping with induced voltages on electronic circuits but the proposed techniques have not been implemented in most of the devices tested in this work. The techniques include the removal of random noise, similar to broadband additive white Gaussian (AWG) noise, or narrow band interferences. In both of those classes of noise, the signals are considered benign. We focused on malicious signal injection using an in-band signal to obscure the actual signal. Disruption to digital circuits by intentional and high-intensity radiation have been investigated by Mansson [4], although signal injection was not directly considered.

Fault injection. The use of hardware fault injection to cause a security breach has been investigated previously [16], [37], [38]. Those attacks are focused only on the digital circuit with the goal of corrupting memory and forcing execution of arbitrary code. This work does not alter the logic of the system, but it exploits a vulnerability in the analog inputs to alter the behavior of the system based on the current unaltered programming logic.

Medical devices. Halperin et al. have demonstrated vulnerabilities specifically in IMDs, but this work is fundamentally different in its approach [39]. Whereas Halperin et al. identified vulnerabilities in the digital control channel used to communicate with IMDs, this work exploits the

analog sensing apparatus of pacemakers and defibrillators. This work also adopts an attack approach that is widely applicable to electronics with analog sensing inputs, not constrained to any particular protocol or specific sensor.

CIEDs have been reported to be affected by static magnetic fields [40] but only at very short distances of under 3 cm, and Low Frequency RFID emissions [24] but the waveforms were not intentionally crafted to force a mis-sense on the device.

In their work on bounding the distance between the communication device and an implantable medical device [41], Rassmussen et al. noted that the receiving microphone was sensitive to electromagnetic interference. Our work goes deeper in analyzing signal injection using RF waveforms. We offer an analysis of the root causes of the vulnerability and develop an attack model using signal injection. We evaluated our attacks on a number of devices including commodity electronics and implantable medical devices.

VII. CONCLUSION

Analog sensors intrinsically trust what they measure, and digital systems trust the input provided by sensors. As a result, intentional electromagnetic interference can trick sensors into providing bogus information to higher-level applications. Implications range from causing pacemakers to stop pacing to injecting chosen touch-tone numbers during phone calls with a Bluetooth headset to an automated bank service center. For distances under 5 cm for an 10 W adversary, our experiments found no clinically relevant risks for completely implanted medical devices. In free air, our experiments caused measurable interference at 1-2 m. We do not believe the current situation reveals an urgent public health risk. Our proposed defenses include traditional analog shielding as well as a digital signal contamination metric based on the root mean square of waveform amplitudes. Our cardiac defense mechanism detects suspicious sensor input by checking whether pacing pulses are consistent with the refractory period of cardiac tissue. Secure websites follow the principle of not trusting unvalidated user input, and the analog of this advice should resonate for sensor systems too.

ACKNOWLEDGMENT

This publication was made possible by Cooperative Agreement No. 90TR0003/01 from the Department of Health and Human Services. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the HHS. This work was also supported by a Sloan Research Fellowship; the University of Minnesota Doctoral Dissertation fellowship; the Korean government (MEST) National Research Foundation (NRF) No. 2012-0000979; the Harvard Catalyst/Harvard Clinical and Translational Science Center MeRIT career development award and the National Science Foundation awards CNS-1035715, CNS-0845671, CNS-0923313, GEO-1124657, and

S12100000211. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Many thanks to Amir Rahmati for creating figures and editing; Quinn Stewart for copy editing, Tingyi Wei for running experiments, the members of the SPQR Lab, Penn's PRECISE center, the UMN Medical Device Center, and the UMN SCLab for feedback on drafts.

REFERENCES

- [1] D. Giri and F. Tesche, "Classification of intentional electromagnetic environments (IEME)," *IEEE Transactions on Electromagnetic Compatibility*, vol. 46, no. 3, pp. 322–328, 2004.
- [2] C. R. Paul, *Electromagnetic Compatibility*. Wiley Online Library, 2005.
- [3] —, "Introduction to electromagnetic compatibility (Wiley Series in Microwave and Optical Engineering)," 2006.
- [4] D. Mansson, R. Thottappillil, and M. Backstrom, "Methodology for classifying facilities with respect to intentional EMI," *IEEE Transactions on Electromagnetic Compatibility*, vol. 51, no. 1, 2009.
- [5] J.-M. Redouté and M. Steyaert, *EMC of analog integrated circuits*. Springer, 2009.
- [6] R. Anderson, *Security Engineering: A guide to building dependable distributed systems*. Wiley, 2010.
- [7] A. Cheng, S. Nazarian, D. D. Spragg, K. Bilchick, H. Tandri, L. Mark, H. Halperin, H. Calkins, R. D. Berger, and C. A. Henrikson, "Effects of surgical and endoscopic electrocautery on modern-day permanent pacemaker and implantable cardioverter-defibrillator systems," *Pacing and clinical electrophysiology*, vol. 31, no. 3, pp. 344–350, 2008.
- [8] J. Loewy, A. Loewy, and E. J. Kendall, "Reconsideration of pacemakers and MR imaging," *Radiographics*, vol. 24, no. 5, pp. 1257–1267, 2004.
- [9] A. Roguin, J. Schwitter, C. Vahlhaus, M. Lombardi, J. Brugada, P. Vardas, A. Auricchio, S. Priori, and T. Sommer, "Magnetic resonance imaging in individuals with cardiovascular implantable electronic devices," *Europace*, vol. 10, no. 3, pp. 336–346, 2008.
- [10] D. Hayes, P. Wang, D. Reynolds, N. Estes, J. Griffith, R. Steffens, G. Carlo, G. Findlay, and C. Johnson, "Interference with cardiac pacemakers by cellular telephones," *New England Journal of Medicine*, vol. 336, no. 21, 1997.
- [11] K. Hekmat, B. Salemin, G. Lauterbach, R. Schwinger, M. Südkamp, H. Weber, and U. Mehlhorn, "Interference by cellular phones with permanent implanted pacemakers: an update," *Europace*, vol. 6, no. 4, pp. 363–369, 2004.
- [12] G. Calcagnini, F. Censi, M. Floris, C. Pignalberi, R. Ricci, G. Biancalana, P. Bartolini, and M. Santini, "Evaluation of electromagnetic interference of GSM mobile phones with pacemakers featuring remote monitoring functions," *Pacing and clinical electrophysiology*, vol. 29, no. 4, 2006.

- [13] W. A. Radasky, C. E. Baum, and M. W. Wik, "Introduction to the special issue on high-power electromagnetics (HPEM) and intentional electromagnetic interference (IEMI)," *IEEE Transactions on Electromagnetic Compatibility*, vol. 46, no. 3, pp. 314–321, 2004.
- [14] M. G. Backstrom and K. G. Lovstrand, "Susceptibility of electronic systems to high-power microwaves: Summary of test experience," *IEEE Transactions on Electromagnetic Compatibility*, vol. 46, no. 3, pp. 396–403, 2004.
- [15] J. Delsing, J. Ekman, J. Johansson, S. Sundberg, M. Backstrom, and T. Nilsson, "Susceptibility of sensor networks to intentional electromagnetic interference," in *17th International Zurich Symposium on Electromagnetic Compatibility*. IEEE, 2006, pp. 172–175.
- [16] S. Govindavajhala and A. Appel, "Using memory errors to attack a virtual machine," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2003.
- [17] American National Standards Institute/Association for the Advancement of Medical Instrumentation (ANSI/AAMI), "Active implantable medical devices — Electromagnetic compatibility — EMC test protocols for implantable cardiac pacemakers and implantable cardioverter defibrillators," 2007.
- [18] W. Irnich, L. Batz, R. Müller, and R. Tobisch, "Electromagnetic interference of pacemakers by mobile phones," *Pacing and clinical electrophysiology*, vol. 19, no. 10, 1996.
- [19] F. Censi, G. Calcagnini, M. Triventi, E. Mattei, and P. Bartolini, "Interference between mobile phones and pacemakers: a look inside," *Ann IST Super Sanita*, vol. 43, no. 3, 2007.
- [20] W. Kainz, J. Casamento, P. Ruggera, D. Chan, and D. Witters, "Implantable cardiac pacemaker electromagnetic compatibility testing in a novel security system simulator," *IEEE Transactions on Biomedical Engineering*, vol. 52, no. 3, 2005.
- [21] H. Moses and J. Mullin, *A practical guide to cardiac pacing*. Lippincott Williams & Wilkins, 2007.
- [22] M. Kroll and M. Lehmann, *Implantable Cardioverter Defibrillator Therapy: The Engineering-Clinical Interface*. Springer, 1996, vol. 188.
- [23] L. Cohan, F. Kusumoto, and N. Goldschlager, "Environmental effects on cardiac pacing systems," *Cardiac Pacing for the Clinician*, 2008.
- [24] S. J. Seidman, R. Brockman, B. M. Lewis, J. Guag, M. J. Shein, W. J. Clement, J. Kippola, D. Digby, C. Barber, and D. Huntwork, "In vitro tests reveal sample radiofrequency identification readers inducing clinically significant electromagnetic interference to implantable pacemakers and implantable cardioverter-defibrillators," *Heart Rhythm*, vol. 7, no. 1, p. 99, 2010.
- [25] C. Sticherling, T. Klingenheben, D. Cameron, and S. H. Hohnloser, "Worldwide clinical experience with a downsized active can implantable cardioverter defibrillator in 162 consecutive patients," *Pacing and clinical electrophysiology*, vol. 21, no. 9, pp. 1778–1783, 2006.
- [26] Syndaver Labs, "Torso #02," http://syndaver.com/product_info.php?cPath=80&products_id=493, Visited Nov 2012.
- [27] D. Hayes, M. Lloyd, and P. Friedman, *Cardiac pacing and defibrillation: a clinical approach*. Wiley-Blackwell, 2000.
- [28] International Electrotechnical Commission, "IEC 60268-16: Sound system equipment-part 16: Objective rating of speech intelligibility by speech transmission index," 2003.
- [29] A. Wang *et al.*, "An industrial strength audio search algorithm," in *Proc. Int. Conf. on Music Info. Retrieval ISMIR*, vol. 3, 2003.
- [30] Weezer, "Island in the sun," Weezer (The Green Album), Geffen Records, Compact Disc, 2001.
- [31] G. A. Miller, "The masking of speech." *Psychological Bulletin*, vol. 44, no. 2, p. 105, 1947.
- [32] Razavi and Behzad, *Design of Analog CMOS Integrated Circuits*, 1st ed. McGraw-Hill, Inc., 2001.
- [33] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms, and Applications*, 4th ed. Pearson Prentice Hall, 2007.
- [34] Z. Jiang, M. Pajic, A. Connolly, S. Dixit, and R. Mangharam, "Real-time heart model for implantable cardiac device validation and verification," in *22nd Euromicro Conference on Real-Time Systems (ECRTS)*, 2010.
- [35] Z. Jiang and R. Mangharam, "Modeling cardiac pacemaker malfunctions with the virtual heart model," in *Annual International Conference on Engineering in Medicine and Biology Society, EMBC*, 2011.
- [36] National Institute of Standards and Technology, "140-2: Security requirements for cryptographic modules," 2001.
- [37] J. Karlsson, P. Folkesson, J. Arlat, Y. Crouzet, G. Leber, and J. Reisinger, "Application of three physical fault injection techniques to the experimental assessment of the MARS architecture," *Dependable Computing and Fault Tolerant Systems*, vol. 10, 1998.
- [38] E. Jenn, J. Arlat, M. Rimen, J. Ohlsson, and J. Karlsson, "Fault injection into VHDL models: the MEFISTO tool," in *Twenty-Fourth International Symposium on Fault-Tolerant Computing (FTCS-24)*, 1994.
- [39] D. Halperin, T. Heydt-Benjamin, B. Ransford, S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno, and W. Maisel, "Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses," in *IEEE Symposium on Security and Privacy*, 2008.
- [40] S. Lee, K. Fu, T. Kohno, B. Ransford, and W. Maisel, "Clinically significant magnetic interference of implanted cardiac devices by portable headphones," *Heart Rhythm*, vol. 6, no. 10, 2009.
- [41] K. Rasmussen, C. Castelluccia, T. Heydt-Benjamin, and S. Capkun, "Proximity-based access control for implantable medical devices," in *Proceedings of the 16th ACM conference on Computer and communications security*, 2009.