Absence Makes the Heart Grow Fonder: New Directions for Implantable Medical Device Security

Tamara Denning[*] *Kevin Fu*[†] *Tadayoshi Kohno*[*]

Abstract

A fundamental tension exists between safety in the common case and security under adversarial conditions for wireless implantable medical devices. We propose a class of new, fail-open defensive techniques for implantable medical devices that attempt to strike a balance between these two goals. We refer to these defensive techniques as *Communication Cloakers*. Cloakers are externally worn devices, much like computational Medical Alert bracelets. Cloakers protect the security of an IMD when worn, but allow for open access during emergencies if removed.

1 Introduction

There is an on-going revolution in wireless *implantable medical device* (*IMD*) technologies. Such devices — which are wholly or partially implanted within patients' bodies — are not only enabling new medical therapies with the potential of greatly improving patients' lives, but are incorporating more sophisticated wireless transceivers and becoming more computationally complex. The latter technological trends, coupled with the physiological importance of these medical devices, suggest potentially harmful consequences if these devices fail to provide appropriate security and privacy safeguards. Indeed, as our own previous research shows [9], it is currently possible for an adversary to use his or her own equipment to reprogram an implantable defibrillator, exploit the defibrillator to compromise the patient's privacy, or even exploit the defibrillator to cause a potentially fatal heart rhythm.

The Challenge. A fundamental challenge is to mitigate the tension between (1) *safety in the common case* and (2) *security under adversarial conditions* for such devices. To understand this tension, consider the simple fact that commercial implantable medical device *programming hardware* (a.k.a. a *commercial programmer*[1]) could be used both by legitimate medical practitioners and by adversaries:

- Under normal, non-adversarial use, the IMD *should* immediately respond to any reprogramming or data read requests from the commercial programmer. This is critical for many implantable medical devices. For example, doctors must wirelessly disable implantable defibrillators before conducting emergency surgeries in order to avoid risk of unnecessary shock to the patients.
- Under adversarial conditions, however, the IMD *should not* respond to reprogramming or data read requests from the commercial programmer since the wireless communications are malicious in intent. It is currently not difficult for adversaries, even without ties to the healthcare community, to obtain a commercial programmer — whether by taking one from a hospital or purchasing one on eBay.

Unfortunately, to the best of our knowledge it is currently difficult or impossible for an IMD to distinguish between these two scenarios. We seek to address this deficiency here, while at the same time also providing security against adversaries capable of developing custom, more sophisticated wireless attack hardware. We believe that our vision, if successful, could significantly improve our ability to provide technical balances between safety in the common case and security under adversarial conditions.

While previous works have directly or indirectly considered this tension, we believe that the previous approaches — while helping mitigate these tensions — are not wholly sufficient. For example, our own earlier recommendation — of making an IMD issue an audible alert during potential adversarial wireless communications [8, 9] — meets our safety goals in the common case since the audible alerts do not impede commercial programmers in clinical settings. However, this approach is also not ideal from a security perspective for some classes of IMDs: a patient might hear an audible alert shortly before a malicious action, but that audible alert will be ineffective if the patient becomes incapacitated shortly thereafter. Additionally, we must consider the potential psychological impact on patients from false-positive alarms. We similarly analyze other existing or conventional approaches for mitigating these tensions in the body of this paper, each of which we argue is insufficient for our purposes.

Failing Open: Communications Cloakers. We step back and propose a class of new directions for balancing safety with security for IMDs. Our directions all in-

[*]Dept. of Computer Science and Engineering, Univ. of Washington.
[†]Dept. of Computer Science, Univ. of Massachusetts Amherst.

[1]Notationally, the medical device community uses the term *programmer* to refer to the external piece of equipment that healthcare practitioners use to wirelessly communicate with and adjust the settings on an implantable medical device.

volve coupling an IMD with an external device — much like the Medical Alert bracelets that IMD patients already wear; however, in contrast to the bracelet, our device is computational in nature and possesses wireless communication capabilities. We use the term *Communication Cloakers* to collectively refer to instances of our approach.

The new twist that all our Cloaker approaches take is to provide security while the patient is wearing a Cloaker, but to provide fail-open access to all external programmers when the patient is not wearing a Cloaker. The model is for the patient to wear a Cloaker during normal, everyday activities (providing security). During normal clinic visits the Cloaker will allow pre-specified, authorized commercial programmers to interact with the IMD, such as the programmer belonging to the prescribing physician's clinic (providing clinical access in the common case). In emergency situations, doctors with previously unauthorized commercial programmers can simply remove the Cloaker from the patient, thereby enabling immediate, emergency open access (providing safety in emergencies).

We believe that the fail-open Cloaker approaches will allow us to achieve new — and we argue, in some cases better — balances between safety and security for many classes of IMDs. Significant details do, however, remain to be fully addressed. We explore the Cloaker design space, the relationship to previous security approaches to IMDs, and some remaining challenges herein. Our study is not definitive, and Cloakers should not be deployed immediately. Rather, in publishing this paper, we seek to encourage broader discussion about the applicability, potential impact, and utility of Cloaker-based approaches for wireless IMD security and safety. We also wish to encourage further study of Cloaker-based and other approaches for mitigating the tensions in wireless IMD design.

2 Design Tensions and Context

Tensions. There are several design tensions that arise from the form factor of IMDs and their role in medical treatment. We have already mentioned the first two above:

- **Safety and Open Access in Emergencies.** An IMD security system needs to allow open access in emergency situations. As noted earlier, emergency caregivers need to have an unhindered ability to change the settings on or disable many classes of IMDs.
- **Security and Privacy Under Adversarial Conditions.** Wireless communication capabilities open up vulnerabilities to attacks from greater distances. A security system for an IMD should strive to meet the traditional security goals of confidentiality, integrity, and availability.
- **Battery Life.** Power is a limited resource for IMDs. IMDs are powered by integrated batteries. Many IMDs have non-rechargeable, non-replaceable batteries, and a drained battery necessitates surgery to implant a replacement. For certain classes of IMDs, such surgeries to replace IMDs bear significant risks and can lead to serious injury and death [6]. It is therefore critical that new security measures are weighed against the effect they will have on the IMD's battery life.
- **Response Time.** Security systems for IMDs should consider the time-sensitivity of the IMD's functionality. Any latencies introduced by security measures should fall within acceptable parameters in order to avoid adversely affecting the patient's health.

The Cloaker approach seeks to balance the first two goals above while simultaneously protecting the battery life and response time of an IMD under both normal and adversarial conditions.

In this paper we consider only attacks instigated by third parties; we do not consider attacks by patients against their own IMDs—for example, patients wishing to illegally self-prescribe new settings on their devices.

On the Importance of Security. It is also worth reflecting upon the likelihood that security for IMDs will become a real problem, or whether it will be a non-issue. While we cannot predict the exact likelihood of wireless computer-based attacks against IMD patients (either directed at specific individuals, or as some form of widespread malice or cyber-terrorism), we can cite specific past examples — all within the past year — of malicious actions against (non-IMD) patients.

The first example, from November 2007, is of malicious attacks against the Coping With Epilepsy website [4]. Attackers placed images on the website that would induce seizures in photosensitive epilepsy patients. This was not an isolated incident, as attackers mounted a similar attack against the Epilepsy Foundation website in March 2008 [15]. These two examples show that malicious parties will attempt to hurt patients via computer-based attacks. While not related to computer-based attacks, our third example suggests that the 81 deaths from Heparin contamination were due to malicious contamination [10].

These past examples suggest that it is important to proactively address the security and privacy of wireless IMDs.

3 Insufficient Approaches

There are a number of possible approaches to IMD communication security. In considering the design space, we

identified several approaches as insufficient for simultaneously addressing both the safety and security of IMDs.

No Security. While one option is to provide no security for an IMD's wireless communications, we find this to be unacceptable. Our earlier work [9] demonstrates that it is possible to compromise the privacy and security of a patient with one model of an implantable defibrillator using a nearby software radio. In Section 2 we also give three examples within the last year in which malcontents appear to have intentionally tried to cause harm — including death — to patients.

Case-by-Case Access Credentials. Any method for obtaining case-by-case access credentials could introduce new failure modes with severe negative safety implications. For example, suppose that an emergency clinician in a foreign country needs to first obtain explicit authorization credentials from a prescribing physician before the clinician can change the settings on or disable an IMD. If the connection between the emergency clinic and the prescribing doctor's office is unavailable, the clinician will not be able to obtain the necessary credentials and therefore may be unable to perform a life-saving operation.

User Alert. In our previous work [8, 9] we suggest having an IMD issue audible alerts or vibrations when it engages in wireless communications. The argument is that these alerts could allow patients to detect the onset of potentially malicious actions, and that these alerts could also serve as a deterrent to adversaries. While we view this approach as raising the bar over the security of existing IMDs such as the defibrillator we studied in [9], we also believe that this approach does not provide an ideal level of security. Indeed, the audible alerts, while informing patients about malicious actions, cannot directly prevent those malicious actions. We reflected on additional drawbacks in Section 1.

Require Close Proximity. Another approach might be to disallow long-distance wireless communications (in meters) until a close-distance bootstrap procedure has completed. For example, an IMD might disallow long-distance communications until after it has verified that the device has the ability to physically contact the patient. Physical-proximity verification can be accomplished by leveraging physiological keying techniques [3], for example, or other methods for signaling that presumably require close or direct contact.

There are a couple of deficiencies with this approach. First, evidence from other technologies suggest that adversaries are often adept at extending the communication ranges of technologies. This means that security should not rely on the presumption that a communications method is short-range-only. See, for example, Kirschenbaum's and Wool's work on building a low-cost RFID skimmer with a longer-than-anticipated communication range [13]; similarly it has been found that certain physiological keying approaches — such as inter-pulse timing [2] — are remotely measurable with vision techniques.

Second, we wish to secure IMDs under threat models in which an attacker can place malicious equipment near a patient. For example, we wish to harden IMD systems against terrorism or other actions involving malicious equipment (placed, for example, near the turnstiles leading into a subway). There are numerous scenarios in which adversarial equipment can be planted near or directly in contact with a patient.

Encryption with Carried Keys. Encrypting communications to and from the IMD can protect the patient from eavesdroppers and unauthorized access. Encryption should be used in the design of a secure system for IMD communications; however, there remains the question of how to distribute keys to legitimate parties.

One seemingly attractive approach would be for patient to carry a card or Medical Alert bracelet imprinted with the IMD's key. There is no guarantee, however, that the card or bracelet will not be lost or damaged in an emergency — or simply forgotten — thereby preventing access to the IMD.

4 New Approaches: Communication Cloakers

We propose a security system in which the presence of a computational device causes the IMD to ignore incoming communications from all other parties and the *absence* of the device causes the IMD to fail open, accepting and responding to all incoming communications. We call the additional computational device a *Communication Cloaker* due to the fact that its presence renders the IMD effectively invisible to unauthorized programmers. In this design the Cloaker mediates communications between the IMD and pre-authorized parties. In emergency situations, the medical staff can remove the Cloaker in order to gain access to the IMD.

Note that this approach allows us to balance safety with security: during day-to-day activities while the patient is wearing his or her Cloaker the patient's IMD is invisible, but during emergency situations — and even when the Cloaker is lost, broken, out of batteries, or simply forgotten — the emergency practitioner can still access the IMD.

We do not seek to dictate exactly how a Cloaker should operate; rather, our goal here is to establish a foundation for exploring the full design space.

System Components. To solidify the above, any Cloaker-based system consists of three components: the

IMD itself; the external programmer; and the Cloaker.

- **IMD.** An IMD is implanted in a patient to treat a medical condition. Examples include implantable pacemakers and defibrillators, drug pumps, and neurostimulators. Modern IMDs can communicate with external devices several meters away.
- **Programmer.** A programmer, or external reader, is used to read data from a patient's IMD and program the device's therapies.
- **Communication Cloaker.** A Communication Cloaker is an additional electronic device that — when worn by the patient — acts as a third-party mediator in the IMD's communications with external programmers.

IMD–Programmer Communications While Wearing the Cloaker. The first challenge we must address is how to manage communications between the IMD and external commercial programmers while the patient is wearing the Cloaker. There are many possibilities, each with their own trade-offs.

We must first consider the communications protocol between the IMD and the Cloaker. Since the Cloaker is the IMD's "buddy" device, they share a long-term relationship and hence can utilize symmetric-only cryptography. The communications must be both encrypted and authenticated, with counters or other sufficient mechanisms to prevent replay and reordering attacks. For our purposes, we assume that the cryptographic layer allows for dropped packets; higher-level protocols can implement a reliable transport on top if necessary.

We must next consider the actual interactions between the IMD, the Cloaker, and the programmer. As one possibility, the IMD could listen for session-initiation requests from the programmer, and then query the Cloaker as an oracle for verifying the authenticity of the programmer. We do not consider this alternative further because it exposes the IMD's battery to a denial-of-service attack (denial-of-sleep attack [17]). Rather, we assume that the Cloaker will first verify that the external programmer is authorized to communicate with the IMD (see below for relevant discussions). Next, we are left with two additional options: whether the Cloaker should proxy the communications between the programmer and the IMD, or whether the Cloaker should hand-off a lightweight (symmetric) access credential to the programmer. The former approach allows the Cloaker to log all the communications between the programmer and the IMD for forensics and analysis purposes, without the potential for packets to be received by the IMD but missed by the Cloaker. The latter approach might provide (potentially negligible) reduction in communications latency, and has the advantage that the Cloaker can subsequently be removed without disrupting the current session.

We return to how the Cloaker verifies that the external programmer is authorized to communicate with the IMD. Here the external nature of the Cloaker — with its replaceable batteries and (potentially) greater computational power — allows significant flexibility. This step can leverage heavyweight public key cryptography. The Cloaker can, for example, be pre-loaded with the public keys of authorized external programmers, such as the programmers for the implanting physician's clinic; the external programmers should, of course, store the corresponding private keys securely in trusted hardware.

IMD–Programmer Communications When Not Wearing the Cloaker. When the patient is not wearing the Communication Cloaker the IMD listens and responds to all incoming communication requests. Unlike hard-coded keys physically printed on Medical Alert bracelets, this approach allows for emergency access even if a patient's Cloaker is lost or damaged as part of an accident: an emergency technician without a pre-authorized external programmer can remove the patient's Cloaker in order to obtain access.

Defining "When Cloaker is Present." A potentially dangerous situation exists if the Cloaker is in the immediate vicinity of a patient — and is therefore inhibiting a programmer's communication with the IMD — but is difficult for emergency medical technicians to locate. For example, the Cloaker could be inside a patient's bag or in a clothes pocket.

One approach for addressing this situation is to incorporate a pulse-sensing unit into the Cloaker, and define the Cloaker as being "present" when it is able to sense a pulse. The Cloaker could be designed to be worn around the patient's wrist, like the Medical Alert bracelets, or around the neck or a finger. The Cloaker could also issue an audible alert when it fails to sense a pulse, thereby warning the patient if the Cloaker accidentally falls off. An adversary could mount a denial-of-service attack under this model in an emergency setting by constantly feeding the Cloaker a pulse even after it is removed from the patient's body, but we view such attacks during emergency settings to be extremely rare, detectable, and counterable by simply destroying the Cloaker.

IMD's Knowledge of the Cloaker's Presence. The next challenge we must address is how to ensure that the IMD knows exactly when the Cloaker is present. Indeed, if we cannot provide such a property, then the attacker might attempt to convince the IMD that the Cloaker is not present when in fact it is, thereby causing the IMD to erroneously enter an open-access state.

As before, there are numerous possibilities, each with their own trade-offs. A stateless approach might involve the IMD querying the Cloaker whenever it detects an ex-

ternal communications request. As a stateful approach, the IMD could instead keep an internal record of the Cloaker's presence and could periodically update this record based on the presence or absence of successful keep-alive messages.

The stateless approach has the advantage of avoiding constant keep-alive messages under non-adversarial circumstances, but as mentioned earlier in a different context, can also expose the IMD to a denial-of-service attack against its battery. Furthermore, suppose the IMD issues an "Are you there?" request immediately after receiving a communication request from a programmer. This predictable behavior might facilitate a malicious programmer's selective jamming of the "Are you there?" messages. Additional logic on the IMD—logic that detects jamming or that sends additional "Are you there?" messages later at non-deterministic times—might be able to mitigate this concern.

There are two possible keep-alive variants. In the first variation, the IMD initiates the keep-alives and the Cloaker sends acknowledgments that the IMD must receive. In the second variation, the Cloaker sends the keep-alive messages. To ensure that the keep-alive messages do not themselves reveal private information to an adversary aside from the obvious fact that wireless communications are taking place, in addition to encrypting and authenticating the messages, we propose using lightweight methods for addressing the wireless packets with non-persistent identifiers [7]. Regardless of which variant is used, the time interval between keep-alive messages in these systems involves a fundamental tradeoff between safety and battery drain; a shorter time interval ensures that the IMD fails open more quickly in the case of an emergency, while a longer time interval requires less power expenditure from the IMD's transceiver.

As with the stateless approach, we must be careful not to allow an attacker to selectively jam the messages from the Cloaker to the IMD. In addition to communicating at non-predictable, pseudorandom time intervals, in-body signaling [19] techniques might also prove applicable. As greater defense-in-depth to signal jamming, we observe that the Cloaker could likely determine that its messages are being jammed (e.g., if it senses a pulse but fails to receive messages or acknowledgments from the IMD). The Cloaker could trigger an alert if it detects such a situation. While in this paper we previously advocated against such alerts as the only defensive mechanisms for an IMD, coupling these alerts with a Cloaker is slightly different since it might allow the patient to reposition himself or herself in order to improve the communications between the Cloaker and the IMD; this proposal does, however, deserve further study. The Cloaker might also increase its transmit power or vary its transmit characteristics upon detecting a jamming attack. As a potentially more controversial recommendation, upon detecting a jamming attack the Cloaker could respond with its own attack to "jam the jammer" — that is, to prevent the jammer from wirelessly communicating with the IMD as it fails open due to the absence of keep-alive messages.

Potential System Extensions. A number of extensions can be made to the core Cloaker system idea in order to enhance its functionality. Recent work by Chae *et al.* [1] suggests that it may be possible to instrument the IMD with zero-power cryptography so that only the Cloaker's encrypted communications cause the IMD's transceiver to wake up. With this cryptography, the system can protect itself against denial-of-sleep attacks on the IMD's keep-alive detection system. The Cloaker system as currently designed assumes that receiving and transmitting messages from the IMD consumes power and therefore seeks to minimize the number of times that the IMD's transceiver must activate. If it is possible to expand the above work such that all IMD transceiver activities are passively powered, then it may be necessary to reexamine the Cloaker system's design.

5 Preliminary Implementation and Evaluation

We developed a proof-of-concept implementation of our proposed system design in order to examine the system's behavior in both standard and adversarial situations and to determine the approximate size of the code base required for the implementation. The system was written in Java in order to maximize the portability of the code base onto different hardware implementations. For our system implementation we chose to keep the IMD aware of the Cloaker's presence via periodic keep-alive messages. In addition, we elected to have the Cloaker forward all authorized communications to the IMD rather than passing off a symmetric session key to the IMD and the authorized programmer.

Our implementation of the Cloaker system accepts input to explicitly affect the state of the system. Input options are able to control system state information such as whether or not the Cloaker is powered and whether or not the Cloaker senses a pulse. In addition, we are able to simulate the effect of a DoS attack on individual system components as well as jamming all wireless communications. The implementation models the communications of one external programmer, the Cloaker, and the IMD. The transceivers of these three components have many functions in common; as a result, the transceivers inherit most of their behavior from a single abstract class. Descriptions of the implementation components are given below. The implementation consists of code that simulates how Cloaker might run on an IMD. Significant fu-

Module Type	Size of Code
Cloaker	179
IMD	115
Programmer	44
Generic Communication/Other	294

Table 1: The total code size divided by module type, calculated as the number of semicolons.

Code Function	Size of Code
I/O	124
Configuration	72
Communication/Functionality	436

Table 2: The total code size divided by code function, calculated as the number of semicolons.

ture work remains to evaluate the proposed approach in the context of a real system.

Cloaker. The Cloaker's functionality is simulated across four different classes controlling the Cloaker's main communication logic, the pulse sensor's state, the timing of the broadcasted keep-alive messages, and the interactive options and command line input.

IMD. The IMD's functionality is modeled by three different classes. These classes control the main logic of the IMD's transceiver, the IMD's responses to authorized commands, and the interactive options and command line input.

Programmer. The external programmer is implemented across two classes that control its main communication logic and its interactive input options.

Wireless Communications. The wireless transmissions are simulated by static methods contained in a single class. We model wireless communications by transferring messages via TCP sockets: when a module's transceiver broadcasts a message the system opens sockets to all other modules and transfers the message.

Code Size. The implementation totals 632 semicolons over 14 classes. Breakdowns of the line counts by module type (Table 1) and by code function (Table 2) are given below. The small size of our implementation indicates that it is feasible to perform line-by-line analysis of the TCB for our IMD security system—something that is highly desirable in a system that will support life-critical functions.

Although we do not have extensive knowledge of the computational capabilities of modern IMDs, we believe that the code base is small enough that it will not tax modern systems. As of 2002, a typical IMD might contain ≥ 1 microprocessors with 32 KHz to 3 MHz clocks and 8 KB to 2 MB of memory (RAM, ROM, EEPROM, Flash). There has been an upward trend in IMD computing resources; pacemakers in the early 1990s typically had ≤ 1 KB of RAM [14].

6 Related Work

There is a growing body of literature on security and privacy in pervasive healthcare. Venkatasubramanian and Gupta [18] provide a survey of current directions within the field, and our own work [8] focuses in particular on implantable medical devices. One previous approach for securing the communications between implantables is to leverage keys derived directly from measurements of the patient's physiology, such as the patient's inter-pulse timing [3].

Our other work highlighted the security issues for IMDs in the context of a real implantable defibrillator, and also proposed new security mechanisms that chipped away at the tension between security and safety [9]. From a defensive direction, we believe that the new Cloaker approaches advocated in this paper will strike a much better balance between security and safety compared to our previous proposals.

The Cloaker approaches have a common ancestor with the RFID Guardian [16], RFID Proxy [12], and the Blocker Tag [11], and indeed all these systems offer similar fail-open behavior; however, whereas the fail-open behavior of these RFID defenses are arguably unintentional side effects of their need for backward-compatibility with existing RFIDs, the fail-open behavior of our Cloakers is explicit and intentional. Having the fail-open characteristic as a design goal, plus the additional unique features available with IMDs, gives us much greater freedom to propose and explore new design trade-offs in the context of IMDs. We can avoid the need to remain backward-compatible with existing IMDs. We can also, for example: directly involve the IMDs in the design of the system; leverage the fact that IMDs are active devices with batteries and greater computational power than RFIDs; and leverage the fact that a Cloaker could have a pulse sensor.

The break-glass system for emergency access to medical records [5] has a model that is related to the Cloaker system; the break-glass system provides one-use passwords for emergency open access to medical records by keeping these passwords in physically tamper-evident containers (e.g., behind a pane of glass).

7 Conclusions and Discussions

We propose a new class of defensive techniques — Communication Cloakers — for improving the security and privacy of wireless implantable medical devices. While Cloakers will not be applicable to all types of IMDs, for many IMDs we argue that Cloakers will strike a new balance between safety in the common case and security un-

der adversarial conditions.

Despite the benefits our new approach, there are still many potential complications worth further investigation. The security of the Cloaker system relies upon the patient's wearing the Cloaker device in any environment where unauthorized communications might take place. If the patient forgets or chooses not to wear the Cloaker device, the security features of the system will be ineffective. Additionally, the action of wearing the Cloaker may have negative cognitive effects on a patient. The Cloaker device itself is more noticeable than many IMDs and may serve as a constant reminder to the patient of his or her condition. Since the device protects against wireless attacks that may or may not occur, the act of wearing the device may cause psychological distress to the patient that is disproportionate to the actual risk involved. We encourage further research on exploring these tensions.

Acknowledgments

We thank William H. Maisel and Benjamin Ransford for their helpful comments on an earlier draft of this paper.

References

[1] H.-J. Chae, D. J. Yeager, J. R. Smith, and K. Fu. Maximalist cryptography and computation on the wisp uhf rfid tag. In *Conference on RFID Security*, July 2007.

[2] S. Chekmenev, A. Farag, and E. Essock. Thermal imaging of the superficial temporal artery: An arterial pulse recovery model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '07)*, pages 1–6, June 2007.

[3] S. Cherukuri, K. Venkatasubramanian, and S. Gupta. BioSec: A biometric based approach for securing communication in wireless networks of biosensors implanted in the human body. In *ICPP Workshops*, 2003.

[4] B. Ertl. Hooligans attack epilepsy patients during epilepsy awareness month. `http://www.pr.com/press-release/60959`, 2007.

[5] A. Ferreira, R. Cruz-Correia, L. Antunes, P. Farinha, E. Oliveira-Palhares, D. W. Chadwick, and A. Costa-Pereira. How to break access control in a controlled manner. In *CBMS '06: Proceedings of the 19th IEEE Symposium on Computer-Based Medical Systems*, pages 847–854, Washington, DC, USA, 2006. IEEE Computer Society.

[6] P. Gould and A. Krahn. Complications associated with implantable cardioverter-defibrillator replacement in response to device advisories. *JAMA*, 295(16):1907–1911, April 2006.

[7] B. Greenstein, D. McCoy, J. Pang, T. Kohno, S. Seshan, and D. Wetherall. Improving wireless privacy with an identifier-free link layer protocol. In *MobiSys*, June 2008.

[8] D. Halperin, T. S. Heydt-Benjamin, K. Fu, T. Kohno, and W. H. Maisel. Security and privacy for implantable medical devices. *IEEE Pervasive Computing*, 7:30–39, January-March 2008.

[9] 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 *IEEE Symposium on Security and Privacy*, May 2008.

[10] G. Harris. Heparin contamination may have been deliberate, F.D.A. says. `http://www.nytimes.com/2008/04/30/health/policy/30heparin.html`, 2008.

[11] A. Juels, R. L. Rivest, and M. Szydlo. The blocker tag: selective blocking of RFID tags for consumer privacy. In *CCS*, 2003.

[12] A. Juels, P. Syverson, and D. Bailey. High-power proxies for enhancing RFID privacy and utility. In *PET*, 2005.

[13] I. Kirschenbaum and A. Wool. How to build a low-cost, extended-range RFID skimmer. In *USENIX Security*, 2006.

[14] W. H. Maisel, W. G. Stevenson, and L. M. Epstein. Changing trends in pacemaker and implantable cardioverter defibrillator generator advisories. *Journal of Pacing and Clinical Electrophysiology*, 25(12):1670–1678, December 2002.

[15] K. Poulsen. Hackers assault epilepsy patients via computer. `http://www.wired.com/politics/security/news/2008/03/epilepsy`, 2008.

[16] M. Rieback, B. Crispo, and A. Tanenbaum. RFID Guardian: A battery-powered mobile device for RFID privacy management. In *ACISP*, pages 184–194, 2005.

[17] F. Stajano and R. J. Anderson. The resurrecting duckling: Security issues for ad-hoc wireless networks. In *Proceedings of Security Protocols, 7th International Workshop, LNCS 1796*, pages 172–194, 1999.

[18] K. Venkatasubramanian and S. Gupta. Chapter 15: Security for pervasive healthcare. In Y. Xiao, editor, *Security in Distributed, Grid, Mobile, and Pervasive Computing*. 2007.

[19] T. Zimmerman. Personal area networks: Near-field intrabody communication. *IBM Systems Journal*, 35(3 & 4):609–617, 1996.