Recent Results in Computer Security for Medical Devices *

Shane S. Clark and Kevin Fu

Department of Computer Science
University of Massachusetts Amherst
{ssclark,kevinfu}@cs.umass.edu

Abstract. The computer security community has recently begun research on the security and privacy issues associated with implantable medical devices and identified both existing flaws and new techniques to improve future devices. This paper surveys some of the recent work from the security community and highlights three of the major factors affecting security and privacy solutions for implantable medical devices: fundamental tensions, software risks, and human factors. We also present two challenges from the security community with which the biomedical community may be able to help: access to medical devices and methods for *in vitro* experimentation.

Key words: security, privacy, imd, wireless, risk, human factors

1 Introduction

The computer security community has shown significantly increased interest in implantable medical devices (IMDs) in the last few years. Security researchers have identified a number of security and privacy flaws in devices that are widely implanted in patients and have begun to suggest technologies to mitigate the associated risks [19, 17, 11, 12]. Many of the issues identified are attributable to the recent widespread adoption of networked, and especially wireless, interfaces in IMDs. Paul Jones of the U.S. Food and Drug Administration has said:

> "The issue of medical device security is in its infancy. This is because, to date, most devices have been isolated from networks and do not interoperate. This paradigm is changing now, creating new challenges in medical device design." (personal communication, Aug. 2007)

IMDs behave like any other networked computing devices in many ways, and, as such, many existing security and privacy risks apply to them. However, computer scientists find IMDs to have unconventional peripherals (e.g., electrical connections to control cardiac tissue). These unique characteristics demand that device designers take care in the adoption of security and privacy mechanisms to this domain.

* This paper appears at the *International ICST Conference on Wireless Mobile Communication and Healthcare* (MobiHealth), October 2011.

Medical Device Security Subtopic	References
Access Control	[5, 8, 9, 10, 11, 19, 20]
Emergent Threats	[12]
Encryption	[5, 10, 11, 17]
Failures	[11, 14, 15, 17, 19]
Foundations & Design Principles	[8, 10, 26]
Hardware	[9, 11, 14, 20]
Human Factors	[4, 23, 26]
Policy	[6, 7, 8]
Privacy	[4, 6, 9, 10, 11, 17, 19, 20]
Software	[7, 8, 12, 15, 18, 26]
Specifications	[6, 7, 8, 18]
Wireless	[5, 8, 9, 10, 11, 17, 19, 20]

Table 1: References to major subtopics in medical device security.

This paper summarizes some of the work done by security researchers in the area of IMDs, with the intention of fostering effective collaborations between the security and biomedical communities. To this end, we discuss three major security and privacy issues that are vital to future research in the area of implantable medical devices: fundamental tensions, software risks, and human factors. We also describe two major challenges for the security community: access to medical devices and methods for *in vitro* experimentation. While not an exhaustive survey, this paper provides the biomedical engineer with several footholds to search for literature about security and privacy for medical devices.

2 Fundamental Tensions: Security, Privacy, Utility, Safety

One of the first key issues with IMDs recognized by security researchers is the existence of fundamental tensions between security and privacy goals and traditional goals such as utility and safety [10]. The strong application of access control and cryptography could endanger patients in the case of an emergency if healthcare professionals are unable to gain access to a device. The status quo, on the other hand, leaves patients vulnerable to malicious parties who could potentially disable an IMD or even use it to induce a life-threatening condition [11, 17, 19]. A balance must be struck between these two extremes.

Security researchers have already proposed some approaches that seek a balance. One proposal is the use of proximity-based access control [20]. By using a technique known as distance bounding, new IMDs could determine how close a device programmer is and only allow access to those nearby. This approach mitigates the effects of malicious behavior because it would require the adversary to be close to the potential victim, potentially alerting the patient to the presence of a threat and allowing him to seek safety. The likelihood that a patient would

recognize an unsafe situation and react appropriately, however, is unknown. The main advantage of an approach like distance bounding is that it would not complicate interactions for medical staff because physical proximity is not a barrier in a clinical setting.

There is no "one size fits all" solution to the tensions between security and privacy and utility and safety. Other proposals include the addition of a second, removable device intended to enforce security and privacy goals [5, 9], a batteryless proxy to handle access control [11], and ultraviolet-ink tattoos to store device keys for emergency access [21]. See Figure 1. Different IMDs bear different risks, so the appropriate balance to strike depends upon the particular class of device under consideration.

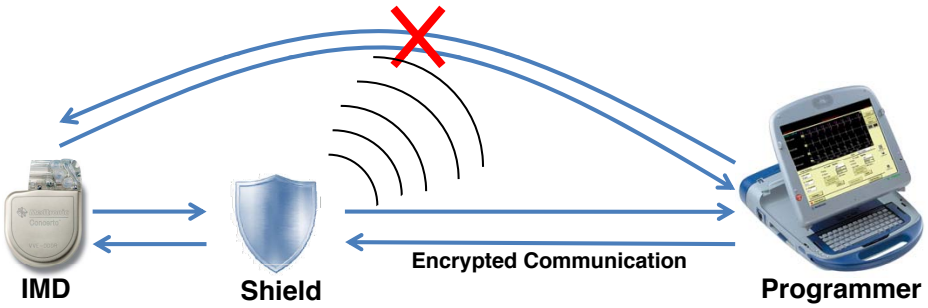

Fig. 1: An illustration of the shield device proposed by Gollakota et al. The shield device jams any radio communication with the IMD unless the programmer has first authenticated with the shield. Removing the shield provides unauthenticated access in case of emergency. Image from [9] used with permission.

3 Hardware vs. Software Security Risks

The increasing size and complexity of the software used in IMDs is another key concern in security and privacy research [13]. As software complexity increases, so do the interactions amongst software artifacts and thus the likelihood of vulnerabilities. For instance, software complexity combined with promiscuous communication can lead to an emergent risk of malware. The recent addition of network interfaces to IMDs massively expands the complexity of the software system that must be protected. Not only must every device designed to communicate with the IMD be trustworthy, but, "any component capable of communication with the device [must] be trustworthy [8]." While it is tempting to write off malware as a concern only for PCs, these devices are increasingly likely to communicate with IMDs and this communication can be used as an infection vector.

There are at least two real-world examples that emphasize the potential impact of malware on networked IMDs. The Stuxnet worm is a computer worm

that propagates through Windows computers and seeks out a specific model of Programmable Logic Controller (PLC) used to control some industrial centrifuges for nuclear enrichment. Once the PLCs are infected, the worm causes the attached centrifuge to spin out of control [3, 24]. The Stuxnet worm is the most widely known example of a computer worm capable of destroying real-world systems, but it is not the only one. More recently, Hanna et al. successfully loaded custom firmware on the Cardiac Science G3 Plus Automated External Defibrillator (AED) because the AED did not verify the authenticity and freshness of a software update. This type of vulnerability is an indication that a malicious party could potentially create a self-replicating worm that spreads to many AEDs. Such a worm could prevent infected AEDs from delivering life-saving shocks or cause them to deliver shocks of arbitrary strength while appearing to function normally [12].

As the complexity and size of a software system increases, the task of ensuring its security and privacy becomes both more important and more difficult to accomplish. There is evidence that the increasing complexity of IMD software is already taking a toll on the biomedical industry. From 1983 to 1997, 6% of the recalls issued for medical devices containing software were attributable to software failures [25]. From 1999 to 2005, 11.3% of the recalls were attributable to software failures, a near doubling of the software-related recall rate [2].

To combat the rising rate of software faults, traditional software-engineering tools such as requirements specification and static analysis should be applied to IMDs [8]. These techniques are used to control complexity and gain confidence in software. Meaningful requirements specification necessitates the consideration of security and privacy at early stages and static analysis tests the final software artifact. An end-to-end approach is essential to developing trustworthy software because it provides confidence in both the design decisions made, and in their correct implementation [16].

4 Security and Human Factors

Security mechanisms must account for human factors. Users must be willing and able to both understand and enforce their own security and privacy goals. If a security mechanism frustrates a user or lacks an intuitive interface, it will serve only to increase the complexity of the system.

The tendency of users to ignore or incorrectly apply security features should not be discounted. In a seminal work in the area of human factors [26], Whitten and Tygar found that only half of the users they tested were able to successfully encrypt an email message—even with access to software manuals and hints from the experimenters. Other work in human factors has demonstrated that users are also likely to ignore security warnings, especially those that become commonplace. Sunshine et al. found that in nine out of their ten user tests, the majority of the test subjects chose to ignore web browser warnings about website certificate validity and proceed to potentially harmful websites [23]. In half of these tests, at least 90% of the subjects chose to ignore the security warnings.

As Whitten and Tygar emphasize, one major obstacle to the adoption and effective use of security and privacy mechanisms is the fact that security is generally a secondary goal. Just as the users in Whitten and Tygar's study simply wanted to send email, IMD patients may simply want to receive treatment for their existing medical conditions. Convincing users to think about and consciously manage a secondary aspect of their IMDs is no small task.

Denning et al. carried out a user study with 13 IMD patients in which the patients were asked about their attitudes toward IMD security and privacy, as well as whether they liked or disliked a variety of security and privacy solutions [4]. The participants' responses revealed that most were concerned about security and privacy in general, but they showed comparatively little concern about specific scenarios. 10 out of 12 participants agreed that they were concerned about the safety and privacy of their electronic information and 9 out of 11 agreed that they were concerned about their physical safety. When asked more detailed questions about security and safety, 10 out of 12 participants disagreed that they were concerned about someone changing their IMD settings without their permission and 7 out of 10 disagreed that they were concerned about medical staff being unable to change settings on their IMDs in the case of an emergency. Despite these somewhat contradictory results, 7 out of 9 participants agreed that something should be done to protect the security of future IMDs [4].

The adoption and usefulness of any new security or privacy technology hinges on user understanding and willingness to participate. Human factors in computer security are still a lively area of research because of the complexities involved with providing understandable options and motivating users to adopt sound security practices. How human factors will affect security and privacy for IMDs is still unclear. More work that specifically addresses this issue is necessary, but any solution that is proposed in the future must take human factors under careful consideration.

5 Challenges for Computer Science

Two major challenges are that (1) computer security researchers seldom have access to real medical devices for experimentation, and (2) the computer security community is largely disjoint from the biomedical engineering community. While security researchers have made some recent progress in understanding the security and privacy issues associated with networked IMDs, the area is still largely unexplored and there are significant barriers to entry.

The ICDs and pacemakers used in all of the published research from the security community were explanted devices obtained from patients or healthcare providers. The explanted devices that our lab uses for experimentation are generally older models, often have aging batteries entering the elective replacement indicator, and none have intact leads. Tests of phenomena such as RF interference are difficult or impossible to carry out in a repeatable way without access to the complete system. It is also difficult to acquire a large number of identical, or even similar, devices.

To make devices more available for reproducible experiments, the Open Medical Device Research Library (OMDRL) now provides explanted medical devices for research in trustworthy computing [1]. Researchers also need open access to hardware-software platforms (e.g., open source pacemakers) to innovate. Otherwise researchers will likely focus on identifying anecdotal vulnerabilities in devices found on eBay rather than innovating new technologies that improve security and privacy.

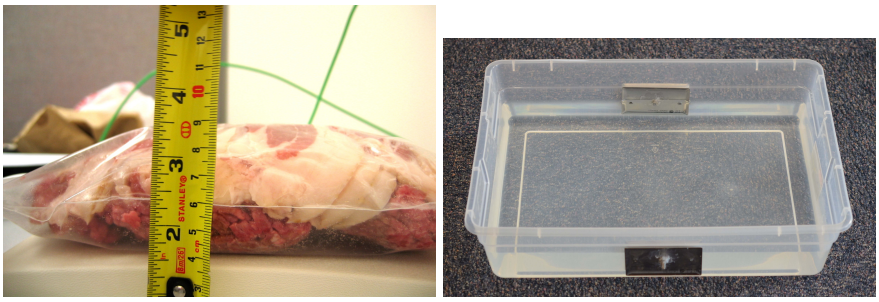

(a) (b)

Fig. 2: (a) The human analogue used by Halperin et al. to prototype defenses [11]. The plastic bag was filled with ground beef and bacon in which a hardware prototype was embedded. (b) Our own FDA-inspired prototype, which is a carefully calibrated saline bath with electrode plates.

The security community also faces major obstacles in designing and completing reproducible *in vitro* experiments because of its isolation from the biomedical community. Security and privacy are issues that must be addressed at a variety of layers. While most of the published vulnerabilities are the result of digital communication interfaces that do not require *in vitro* experiments to validate, many of the proposed defenses rely on physical-layer properties [11, 20, 9]. Halperin et al. were the first security researchers to attempt a realistic *in vitro* experiment. They chose to use a plastic bag full of hamburger and bacon to approximate a human torso (see Figure 2) [11]. Despite a lack of scientific rigor, bags of meat remain the state of the art testing methodology for computer scientists working with IMDs. Lacking familiarity with the literature from the biomedical community, it is difficult for security researchers to determine proper procedures, or even to find the appropriate standards to follow. Our own group has recently begun prototyping new testing setups based on published literature from the biomedical community (see Figure 2) [22]. The state of *in vitro* testing in the security community is improving, but the biomedical community still has an opportunity to improve the testing methodology used by actively engaging in collaborations with security researchers.

6 Conclusion

This paper presents three major considerations that must be addressed by researchers working on security and privacy for IMDs and outlines two challenges from the computer science community. Our hope is that future research can leverage the strengths of both the computer science and biomedical communities to produce new and effective approaches to IMD security and privacy.

Acknowledgments

We thank Wayne Burleson, Gesine Hinterwalder, and Ben Ransford for their feedback on early drafts. This research is supported by NSF CNS-0831244, a Sloan Research Fellowship, an NSF graduate research fellowship, and 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 DHHS or NSF.

References

1. Open Medical Device Research Library, http://www.omdrl.org/
2. Bliznakov, Z., Mitalas, G., Pallikarakis N.: Analysis and Classification of Medical Device Recalls. World Congress on Medical Physics and Biomedical Engineering (2006)
3. Broad, W.J., Markoff, J., Sanger, D.E.: Israeli Test on Worm Called Crucial in Iran Nuclear Delay. http://www.nytimes.com/2011/01/16/world/middleeast/16stuxnet.html (2011)
4. Denning, T., Borning, A., Friedman, B., Gill, B.T., Kohno, T., Maisel, W.H.: Patients, Pacemakers, and Implantable Defibrillators: Human Values and Security for Wireless Implantable Medical Devices. International Conference on Human Factors in Computing Systems (2010)
5. Denning, T., Fu, K., Kohno, T.: Absence Makes the Heart Grow Fonder: New Directions for Implantable Medical Device Security. USENIX Workshop on Hot Topics in Security (2008)
6. Fu, K.: Inside Risks, Reducing the Risks of Implantable Medical Devices: A Prescription to Improve Security and Privacy of Pervasive Health Care. Communications of the ACM, vol. 52(6), pp. 25–27 (2009)
7. Fu, K.: Software Issues for the Medical Device Approval Process. Statement to the Special Committee on Aging, United States Senate, Hearing on a Delicate Balance: FDA and the Reform of the Medical Device Approval Process (2011)
8. Fu, K.: Trustworthy Medical Device Software. Public Health Effectiveness of the FDA 510(k) Clearance Process: Measuring Postmarket Performance and Other Select Topics, Workshop Report (2011)
9. Gollakota, S., Hassanieh, H., Ransford, B., Katabi, D., Fu, K.: They Can Hear Your Heartbeats: Non-Invasive Security for Implanted Medical Devices. ACM SIGCOMM (2011)
10. Halperin, D., Heydt-Benjamin, T.S., Fu, K., Kohno, T., Maisel, W.H.: Security and Privacy for Implantable Medical Devices. IEEE Pervasive Computing, vol. 7, pp. 30–39 (2008)

11. Halperin, D., Heydt-Benjamin, T.S., Ransford, B., Clark, S.S., Defend, B., Morgan, W., Fu, K., Kohno, T., Maisel, W.H.: Pacemakers and Implantable Cardiac Defibrillators: Software Radio Attacks and Zero-Power Defenses. IEEE Symposium on Security and Privacy (2008)
12. Hanna, S., Rolles, R., Molina-Markham, A., Fu, K., Song, D.: Take Two Software Updates and See Me in the Morning: The Case for Software Security Evaluations of Medical Devices. USENIX Workshop on Health Security and Privacy (2011)
13. Israel, C.W., Barold, S.S.: Pacemaker Systems as Implantable Cardiac Rhythm Monitors. American Journal of Cardiology (2001)
14. Lee, S., Fu, K., Kohno, T., Ransford, B., Maisel, W.H.: Clinically Significant Magnetic Interference of Implanted Cardiac Devices by Portable Headphones. Heart Rhythm Journal vol. 6(10), pp. 1432–1436 (2009)
15. Leveson, N.G., Turner, C.S.: An Investigation of the Therac-25 Accidents. Computer, vol. 26(7), pp. 18–41 (1993)
16. Leveson, N.G.: Safeware: System Safety and Computers. Addison-Wesley (1995)
17. Li, C., Raghunathan, A., Jha, N.K.: Hijacking an Insulin Pump: Security Attacks and Defenses for a Diabetes Therapy System. IEEE International Conference on e-Health Networking, Applications, and Services (2011)
18. Networking and Information Technology Research and Development Program: High-Confidence Medical Devices: Cyber-Physical Systems for 21st Century Health Care (2009)
19. Paul, N., Klonoff, D.C.: Insulin Pump System Security and Privacy. USENIX Workshop on Health Security and Privacy (2010)
20. Rasmussen, K.B., Castelluccia, C., Heydt-Benjamin, T.S., Capkun, S.: Proximity-Based Access Control for Implantable Medical Devices. ACM Conference on Computer and Communications Security (2009)
21. Schechter, S.: Security that is Meant to be Skin Deep: Using Ultraviolet Micropigmentation to Store Emergency-Access Keys for Implantable Medical Devices. USENIX Workshop on Health Security and Privacy (2010)
22. Seidman, S.J., Ruggera, P.S., Brockman, R.G., Lewis, B., Shein, M.J.: Electromagnetic Compatibility of Pacemakers and Implantable Cardiac Defibrillators Exposed to RFID Readers. International Journal on Radio Frequency Identification Technology and Applications (2007)
23. Sunshine, J., Egelman, S., Almuhimedi, H., Atri, N., Cranor, L.F.: Crying Wolf: An Empirical Study of SSL Warning Effectiveness. USENIX Security Symposium (2009)
24. The Stuxnet Worm, http://www.symantec.com/business/outbreak/index.jsp?id=stuxnet/
25. Wallace, D., Kuhn, D.: Failure Modes in Medical Device Software: An Analysis of 15 Years of Recall Data. International Journal of Reliability Quality and Safety Engineering (2001)
26. Whitten, A., Tygar, J.: Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. USENIX Security Symposium (1999)