IEEE pervasive COMPUTING
MOBILE AND UBIQUITOUS SYSTEMS
www.computer.org/pervasive

Security and Privacy for Implantable Medical Devices

Daniel Halperin, Thomas S. Heydt-Benjamin, Kevin Fu, Tadayoshi Kohno, and William H. Maisel

Vol. 7, No. 1
January–March 2008

Security and Privacy for Implantable Medical Devices

Protecting implantable medical devices against attack without compromising patient health requires balancing security and privacy goals with traditional goals such as safety and utility.

Daniel Halperin and Tadayoshi Kohno
University of Washington

Thomas S. Heydt-Benjamin and Kevin Fu
University of Massachusetts Amherst

William H. Maisel
Beth Israel Deaconess Medical Center and Harvard Medical School

Implantable medical devices monitor and treat physiological conditions within the body. These devices—including pacemakers, implantable cardiac defibrillators (ICDs), drug delivery systems, and neurostimulators—can help manage a broad range of ailments, such as cardiac arrhythmia, diabetes, and Parkinson's disease (see the "Pacemakers and Implantable Cardiac Defibrillators" sidebar). IMDs' pervasiveness continues to swell, with upward of 25 million US citizens currently reliant on them for life-critical functions.[1] Growth is spurred by geriatric care of the aging baby-boomer generation, and new therapies continually emerge for chronic conditions ranging from pediatric type 1 diabetes to anorgasmia and other sexual dysfunctions. Moreover, the latest IMDs support delivery of telemetry for remote monitoring over long-range, high-bandwidth wireless links, and emerging devices will communicate with other interoperating IMDs.

Despite these advances in IMD technologies, our understanding of how device security and privacy interact with and affect medical safety and treatment efficacy is still limited. Established methods for providing safety and preventing unintentional accidents (such as ID numbers and redundancy) don't prevent intentional failures and other security and privacy problems (such as replay attacks). Balancing security and privacy with safety and efficacy will become increasingly important as IMD technologies evolve. To quote Paul Jones from the US Food and Drug Administration, "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).

We present a general framework for evaluating the security and privacy of next-generation wireless IMDs. Whereas others have considered specific mechanisms for improving device security and privacy, such as the use of physiological values as encryption keys for inter-IMD communication (see the "Related Work in Implantable-Medical-Device Security" sidebar),[2] we ask a broader question: What should be the security and privacy design goals for IMDs? When we evaluate these goals in the broader context of practical, clinical deployment scenarios, we find inherent tensions between them and traditional goals such as safety and utility. To further complicate matters, the balance between security, privacy, safety, and utility might differ depending on the IMD in question. We also present a set of possible research directions for mitigating these tensions. Our framework and follow-on research will help provide a foundation for IMD manufacturers—as well as regulatory bodies such as the FDA—to evaluate, understand, and

Published by the IEEE CS ■ 1536-1268/08/$25.00 © 2008 IEEE

address the security and privacy challenges created by next-generation wireless IMDs.

Criteria for implantable medical devices

We suggest several candidate criteria for IMDs. A particular criterion's applicability might vary, depending on the type of IMD.

Safety and utility goals

Traditional IMD design goals include safety—the IMD should net much greater good than harm—and utility—the IMD should be useful to both clinicians and patients. For our purposes, these goals encompass other goals, such as reliability and treatment efficacy. Our survey of utility and safety goals for IMDs focuses on those that potentially conflict with IMD security and privacy.

Data access. Data should be available to appropriate entities. For example, many devices must report measured data to healthcare professionals or certain physiological values to patients.

In emergency situations, IMDs can provide useful information to medical professionals when other records might be unavailable. Many existing devices present information such as a patient's name and sometimes a stored diagnosis and history of treatments. They could also contain medical characteristics such as allergies and medications.

Data accuracy. Measured and stored data should be accurate. For patient monitoring and treatment, this data includes not only measurements of physiological events but also a notion of when those events occurred.

Device identification. An IMD should make its presence and type known to authorized entities. A caregiver frequently needs to be aware of an IMD's presence. For example, an ICD should be deactivated before surgery. For this reason, the FDA recently considered attaching remotely readable RFID tags to implanted devices.[3]

Configurability. Authorized entities should be able to change appropriate IMD settings. For example, doctors should be able to choose which therapies an ICD will deliver, and patients with devices such as open-loop insulin pumps need partial control over the settings.

Updatable software. Authorized entities should be able to upgrade IMD firmware and applications. Appropriately engineered updates can be the safest way to recall certain classes of IMDs because the physical explantation of some devices—such as pacemakers and ICDs—can lead to serious infection and death.

Multidevice coordination. Although some examples of inter-IMD communications exist (such as contralateral routing of signal [CROS] hearing aids), projected future IMD uses involve more advanced coordinated activities.[4] For example, a future closed-loop insulin delivery system might automatically adjust an implanted insulin pump's settings on the basis of a continuous glucose monitor's readings.

Auditable. In the event of a failure, the manufacturer should be able to audit the device's operational history. The data necessary for the audit might differ from the data exposed to healthcare professionals and patients via typical data access.

Resource efficient. To maximize device lifetime, IMDs should minimize power consumption. Newer IMDs enhanced with wireless communications will expend more energy than their passive predecessors, so they must minimize computation and communication. IMD software should also minimize data storage requirements.

Security and privacy goals

To understand the unique challenges of balancing security and privacy with safety and effectiveness, we first review how the standard principles of computer security—including confidentiality, integrity, and availability—extend to IMDs. We focus on security and privacy goals for IMDs themselves, deferring to other works for a discussion of how to protect a patient's IMD data after it's stored on a back-end server (see the "Related Work in Implantable-Medical-Device Security" sidebar).

In the event of a failure, the manufacturer should be able to audit the implantable medical device's operational history.

Authorization. Many goals of secure IMD design revolve around authorization, which has several broad categories:

- *Personal authorization.* Specific sets of people can perform specific tasks. For example, patients or primary-care physicians might be granted specific rights after authentication of their personal identities. Depending on the authentication scheme, these rights might be

Pacemakers and Implantable Cardiac Defibrillators

Both pacemakers and ICDs are designed to treat abnormal heart conditions. About the size of a pager, each device is connected to the heart via electrodes and continuously monitors the heart rhythm. Pacemakers automatically deliver low-energy signals to the heart to cause the heart to beat when the heart rate slows. Modern ICDs include pacemaker functions but can also deliver high-voltage therapy to the heart muscle to shock dangerously fast heart rhythms back to normal. Pacemakers and ICDs have saved innumerable lives,[1] and the number of ICD implants will soon exceed 250,000 annually.[2]

Internals

Pacemakers and ICDs typically consist of a sealed, battery-powered, sensor-laden pulse generator; several steroid-tipped, wire electrodes (leads) that connect the generator to the myocardium (heart muscle); and a custom ultralow-power microprocessor, typically with about 128 Kbytes of RAM for telemetry storage.[3] The device's primary function is to sense cardiac events, execute therapies, and store measurements such as electrocardiograms. Healthcare professionals configure the settings on pacemakers and ICDs using an external device called a programmer.

Pacemakers and ICDs often contain high-capacity lithium-based batteries that last five to seven years.[4] Rechargeable batteries are extremely rare, for practical, economic, and safety reasons. Device lifetime depends on the treatments required. Whereas pacing pulses consume only about 25 µJ, each ICD shock consumes 14 to 40 J.[4] A single defibrillation can reduce the ICD's lifetime by weeks.

Wireless communications

Previous generations of pacemakers and ICDs communicated at low frequencies (near 175 kHz) with a short read range (8 cm) and used low-bandwidth (50 Kbits per second) inductive coupling to relay telemetry and modify therapies.[5] Modern devices use the Medical Implant Communications Service, which operates in the 402- to 405-MHz band and allows for much higher bandwidth (250 Kbps) and longer read range (specified at two

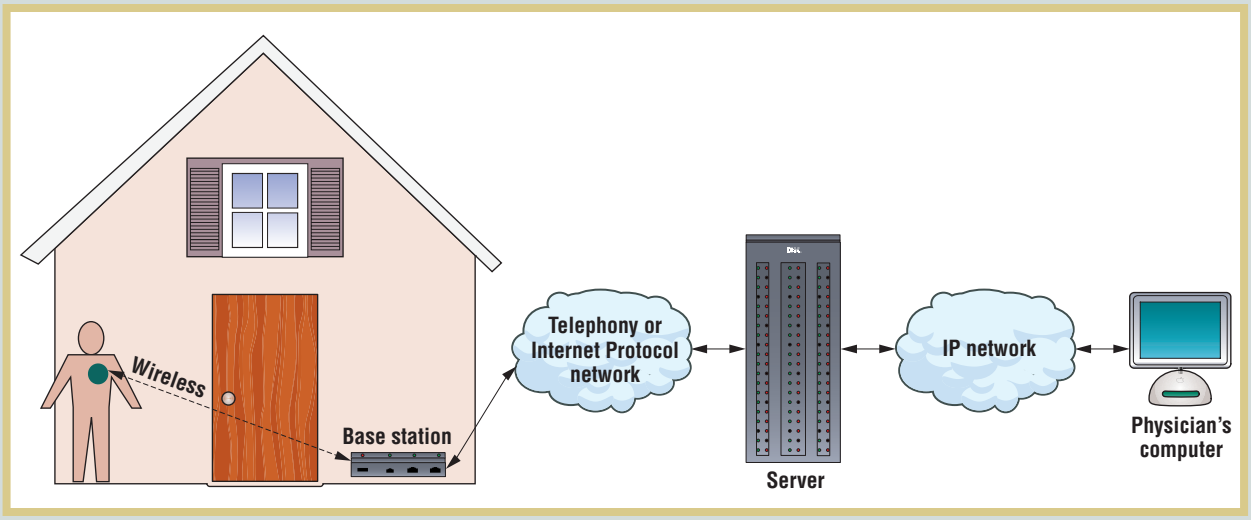

Figure A. Recent implantable cardiac defibrillators provide home monitoring via wireless base stations that relay data to doctors with Web access.

delegatable to other entities.

- *Role-based authorization*. An entity is authorized for a set of tasks on the basis of its role, such as physician or ambulance computer. The device manufacturer might also have special role-based access to the device.
- *IMD selection*. When an external entity communicates with one or more IMDs, it must ensure it communicates with only the intended devices.

Authorization and authentication in a medical setting can be highly sensitive to context. For example, a device might be configured to relax authorization rules if it detects an emergency condition, under the assumption that the patient will suffer less harm from weakly authorized (or even anonymous) intervention than from no intervention. Such context awareness for IMDs is related to the criticality-aware access-control model.[5] Regardless of policy, an IMD should

to five meters).[5] As figure A illustrates, major pacemaker and ICD manufacturers now produce at-home monitors that wirelessly collect data from implanted devices and relay it to a central repository over a dialup connection. The repository is accessible to doctors via an SSL-protected Web site.

Reliability

Although pacemakers and ICDs often save lives, they can occasionally malfunction. Safety issues involving these devices have received much attention. Since 1990 the US Food and Drug Administration has issued dozens of product advisories affecting hundreds of thousands of pacemakers and ICDs.[6] These statistics show that 41 percent of device recalls were due to malfunctions in firmware (216,533 out of 523,145 devices). Additional device programming glitches remain, as evidenced by a clock function abnormality that we recently observed in a clinical setting (see figure B).

These problems' existence underscores potential hazards that come with increasingly sophisticated implantable medical devices. Past abnormalities surfaced under accidental circumstances. The potential for intentionally malicious behavior calls for a deeper investigation into IMD safety from a security and privacy perspective.

REFERENCES

1. W. Maisel, "Safety Issues Involving Medical Devices," *J. Am. Medical Assoc.*, vol. 294, no. 8, 2005, pp. 955–958.
2. M. Carlson et al., "Recommendations from the Heart Rhythm Society Task Force on Device Performance Policies and Guidelines," *Heart Rhythm*, vol. 3, no. 10, 2006, pp. 1250–1273.
3. C. Israel and S. Barold, "Pacemaker Systems as Implantable Cardiac Rhythm Monitors," *Am. J. Cardiology*, Aug. 2001, pp. 442–445.
4. J. Webster, ed., *Design of Cardiac Pacemakers*, IEEE Press, 1995.
5. H. Savci et al., "MICS Transceivers: Regulatory Standards and Applications [Medical Implant Communications Service]," *Proc. IEEE SoutheastCon 2005*, IEEE Press, 2005, pp. 179–182.
6. W. Maisel et al., "Recalls and Safety Alerts Involving Pacemakers and Implantable Cardioverter-Defibrillator Generators," *J. Am. Medical Assoc.*, vol. 286, no. 7, 2001, pp. 793–799.

Arrhythmia Logbook Report

Episode Query Selections

Show All Episodes

Episode	Date/Time	Type	Zone	Rate bpm	Therapy/ Duration
1,230	23-JUN-05 19:10	Spont	VF	222	Diverted
1,229	20-JUN-05 12:08	Spont	VF	216	Diverted
1,228	21-MAY-07 21:22	ATR		130	??:??
1,227	21-MAY-07 15:01	ATR		121	06:20 h:m
1,226	21-MAY-07 15:01	ATR		119	00:45 m:s
1,225	21-MAY-07 15:00	ATR		120	00:11 m:s
1,224	21-MAY-07 15:00	ATR		119	00:16 m:s
1,223	21-MAY-07 15:00	ATR		118	00:07 m:s
1,222	21-MAY-07 14:59	ATR		119	00:09 m:s

Figure B. A report from a patient's routine ICD check. Throughout the device's lifetime, 1,230 arrhythmia episodes have occurred and been automatically recorded (column 1). Episodes with higher numbers occur after episodes with lower numbers. Yet, the ICD incorrectly notes the date and time for episodes 1,229 and 1,230, reporting them as occurring in 2005 when they actually occurred in 2007.

have the technological means to enforce the authorization goals.

Availability. An adversary should not be able to mount a successful denial-of-service (DoS) attack against an IMD. For example, an adversary should not be able to drain a device's battery, overflow its internal data storage media, or jam any IMD communications channel.

Device software and settings. Only authorized parties should be allowed to modify an IMD or to otherwise trigger specific device behavior (for example, an outsider should not be able to trigger an ICD's test mode, which could induce heart failure). Physicians or device manufacturers should place bounds on the settings available to patients to prevent them from accidentally or intentionally harming themselves (for instance, patients should not be able to increase morphine delivery from an implanted

Related Work in Implantable-Medical-Device Security

Much research focuses on securing computer-based medical devices against unintentional failures, such as accidents in radiation treatments from the Therac-25.[1] Interest in protecting these devices against intentional failures is increasing. In a survey of current security directions in pervasive healthcare, Krishna Venkatasubramanian and Sandeep Gupta focus on these aspects:[2]

- efficient methods for securely communicating with medical sensors, including IMDs (such as BioSec's use of physiological values as cryptographic keys);
- controlling access to patient data after aggregation into a management plane (Marci Meingast, Tanya Roosta, and Shankar Sastry provide another discussion[3]); and
- legislative approaches for improving security.

Although others consider the security and privacy of IMD data management by external applications, our research focuses on the challenges and design criteria inherent in IMDs themselves. Even when focusing solely on IMDs, we find fundamental tensions between the security, privacy, safety, and utility goals—particularly when evaluating these goals in the broader context of realistic usage scenarios. Simply using secure communications protocols can't solve these tensions. Finding a suitable balance between these tensions is nontrivial. John Halamka and his colleagues began this process in the context of the VeriChip RFID tag, a low-end implantable device.[4] They concluded that the VeriChip tag shouldn't be used for certain security-sensitive purposes. Jason Hong and his colleagues consider models for tackling the problem of balanced privacy for ubiquitous computing systems.[5]

Although many of the issues we raise are applicable to non-IMD medical devices, IMDs have unique characteristics. For example, replacing certain IMDs through surgery can be risky, and even deadly,[6] so certain IMDs should have long battery lives or be remotely rechargeable. Additionally, unlike other medical devices, IMDs are designed to be part of a patient's everyday, nonclinical activities, thus increasing the opportunity for security or privacy violations.

REFERENCES

1. N.G. Leveson and C.S. Turner, "An Investigation of the Therac-25 Accidents," *Computer*, vol. 26, no. 7, 1993, pp. 18–41.
2. K. Venkatasubramanian and S. Gupta, "Security for Pervasive Healthcare," *Security in Distributed, Grid, Mobile, and Pervasive Computing*, Y. Xiao, ed., CRC Press, 2007, pp. 349–366.
3. M. Meingast, T. Roosta, and S. Sastry, "Security and Privacy Issues with Health Care Information Technology," *Proc. Int'l Conf. Eng. Medicine and Biology Soc.* (EMBS 06), IEEE Press, 2006, pp. 5453–5458.
4. J. Halamka et al., "The Security Implications of VeriChip Cloning," *J. Am. Medical Informatics Assoc.*, vol. 13, no. 6, 2006, pp. 601–607.
5. J. Hong et al., "Privacy Risk Models for Designing Privacy-Sensitive Ubiquitous Computing Systems," *Proc. 5th Conf. Designing Interactive Systems* (DIS 04), ACM Press, 2004, pp. 91–100.
6. P. Gould and A. Krahn, "Complications Associated with Implantable Cardioverter-Defibrillator Replacement in Response to Device Advisories," *J. Am. Medical Assoc.*, vol. 295, no. 16, 2006, pp. 1907–1911.

pump). Similarly, the physician can have access to modify most device settings but should not have unrestricted access to the audit logs or debug modes. IMDs should only accept authorized firmware updates.

Device-existence privacy. An unauthorized party should not be able to remotely determine that a patient has one or more IMDs. An adversary might be a potential employer willing to discriminate against the ill, a member of an organized-crime group seeking to sell a valuable device, or, in the case of military personnel, an enemy operative.

Device-type privacy. If a device reveals its existence, its type should still only be disclosed to authorized entities. Patients might not wish to broadcast that they have a particular device for many reasons. For example, the device might treat a condition with a social stigma, it might be associated with a terminal condition, or it might be extremely expensive.

Specific-device ID privacy. An adversary should not be able to wirelessly track individual IMDs. This is analogous to the concern about the use of persistent identifiers in RFIDs,[6] Bluetooth,[7] and 802.11 media access control (MAC) addresses[8] to compromise an individual's location privacy.

Measurement and log privacy. Consistent with standard medical privacy practices, an unauthorized party should not be able to learn private information about the measurements or audit log data stored on the device. The adversary should also not be able to learn private information about ongoing telemetry.

Bearer privacy. An adversary should not be able to exploit an IMD's properties to identify the bearer or extract private (nonmeasurement) information about the patient. This information includes a patient's name, medical history, or detailed diagnoses.

Data integrity. An adversary should

not be able to tamper with past device measurements or log files or induce specious modifications into future data. No one should be able to change when an event occurred, modify its physiological properties, or delete old events and insert new ones. A patient's name, diagnoses, and other stored data should be tamper-proof.

Classes of adversaries

No treatment of security is complete without a discussion of adversarial resources. For our purposes, the set of adversaries includes, but isn't limited to, these:

- *Passive adversaries*. Such adversaries eavesdrop on signals (both intentional and side-channel) transmitted by the IMD and by other entities communicating with the IMD.
- *Active adversaries*. These adversaries can also interfere with legitimate communications and initiate malicious communications with IMDs and external equipment.
- *Coordinated adversaries*. Two or more adversaries might coordinate their activities—for example, one adversary would be near a patient and another near a legitimate IMD programmer.
- *Insiders*. Insiders can be potential adversaries. Examples include healthcare professionals, software developers, hardware engineers, and, in some cases, patients themselves.

We further subdivide each of these categories by the equipment the adversaries use:

- *Standard equipment*. Adversaries might use commercial equipment for malicious purposes. For instance, they might steal a device programmer from a clinic.
- *Custom equipment*. Adversaries might develop home-brewed equipment for eavesdropping or active attacks. This equipment could have additional amplification, filtering, and directional antennas, and isn't limited to legal bounds on transmitter power or other parameters.

Tensions

As we mentioned earlier, inherent tensions exist between some security and privacy goals and traditional goals such as utility and safety.

Security versus accessibility

Consider two scenarios.

In the first scenario, an unconscious patient with one or more IMDs enters an emergency room, perhaps in a foreign country or developing region. Emergency-room personnel quickly determine the types of IMDs the patient has. The staff then use standard equipment to interrogate the IMDs, extract critical physiological information, and treat the patient, including altering IMD settings and even firmware as appropriate. Because the patient is alone and has no form of identification, the staff also extracts the patient's name and other pertinent information from the data stored on the IMDs.

This scenario corresponds to the current technology in deployed IMDs and external programmers.

In the second scenario, a patient explicitly controls which individuals—or specific external devices—can interact with his or her IMDs. The IMDs use strong access-control and cryptographic mechanisms to prevent unauthorized exposure of data and unauthorized changes to settings. The IMDs also use mechanisms to provide bearer, specific-device ID, device-type, and device-existence privacy.

In this scenario, most of our security criteria are met.

Notice how these two scenarios are diametrically opposed. If a patient's IMDs use strong security mechanisms, as outlined in the second scenario, the equipment in an unfamiliar emergency room won't be authorized to discover, access, or otherwise interact with the patient's IMDs. The emergency-room technicians wouldn't have access to information about the patient's physiological state immediately before his or her admittance to the hospital. Without knowledge of IMD existence, administered care could be dangerous, and the inability to alter settings or deactivate IMDs might prevent necessary treatment because some IMDs (such as ICDs) might need to be deactivated to avoid risk of injury to the surgeons.

The most natural approach for providing access to an IMD in emergency situations would be to incorporate back doors for emergency-room equipment. However, an adversary could exploit the back doors.

> The use of cryptography can create tension between security and some implantable medical devices' longevity and performance goals.

Security versus device resources

Strong security mechanisms, such as public-key cryptography, can be expensive in terms of both computational time and energy consumption. As with general sensor networks, the use of cryptography can therefore create tension between security and some IMDs'

longevity and performance goals. Moreover, increasing resource use for secure communications can amplify the effects of certain malicious DoS attacks, such as repeated attempts to authenticate.

For security, IMDs might also wish to keep detailed records of all transactions with external devices (we elaborate on this later). These transaction logs could potentially overflow a device's onboard memory, particularly under DoS attacks or when an adversary explicitly seeks to exhaust a device's memory.

Security versus usability

The standard tension between security and usability also applies to IMDs. From a usability perspective, long-distance wireless communication between IMDs and external devices offers many advantages, including continuous at-home monitoring and flexibility in clinical settings. But, from a security perspective, wider-range wireless communications increases exposure to both passive and active adversaries. In the Medical Implant Communications Service (MICS) band, an attacker with limited resources might extend the specification's five-meter distance using a directional antenna and an inexpensive amplifier.

Furthermore, the careful addition of new security mechanisms shouldn't overly complicate user interfaces on the external devices, particularly when healthcare professionals must make quick decisions during emergency care.

Research directions

Although completely eliminating tensions between the various goals for IMDs might be impossible, several directions deserve further research and exploration. We confine ourselves primarily to a high-level examination of these directions, some of which we plan to build on in future research. We focus on security- and privacy-related research; other advances in technology (such as longer battery lives or safer methods for device replacement) might also mitigate some tensions.

Fine-grained access control

The two scenarios we described demonstrate a tension between open access to devices during emergency situations and the use of prespecified access-control lists. In the first scenario, emergency caregivers will be able to communicate with a patient's IMD, but so will an adversary. Conversely, the latter scenario could prevent adversarial access to an IMD but will also lock out an emergency caregiver.

If we assume that emergency technicians' external programmers will always be connected to the Internet, easing the tension between these two goals might be possible. The programmer could first interrogate the patient's IMD to learn the device's manufacturer, model, serial number, and possibly the patient's primary-care facility. The programmer could then contact the manufacturer or primary-care facility. The manufacturer or primary-care facility could review the request and, much like the Grey system,[9] issue a signed credential granting the programmer the rights to access specific IMD functions for a specified time period.

This approach would help ensure that the manufacturer or primary-care facility has ultimate control over which external devices can interact with a particular IMD. However, this approach isn't conducive to specific-device ID privacy. It might also introduce safety concerns if the Internet connection between the emergency technician's programmer and the device manufacturer or primary-care facility is slow, severed, or otherwise faulty.

Open access with revocation and second-factor authentication

The medical community might decide that it's sufficient to always allow commercial medical equipment to access IMDs if it is possible to revoke or limit access from lost or stolen equipment. For example, revocation could occur implicitly through automatically expiring certificates for IMD programmers. These certificates should be hard to re-obtain without proper medical credentials, should be stored in secure hardware, and might be distributed hierarchically in a clinic. However, this approach exposes IMDs to compromised equipment for short periods, requires them to have a secure and robust notion of time, and opens a caregiver to potential DoS attacks through the certificate distribution system. The requirement to coordinate the use of such an infrastructure across international boundaries for it to function on a global scale might limit its potential.

IMD programmers could also require a secondary authentication token, such as a smart card, tied to a medical professional's identity. Requiring such tokens could further limit unauthorized parties' use of legitimate medical equipment, although it might also decrease usability and increase emergency response time. Alternatively, manufacturers might be able to extend sophisticated federated iden-

From a security perspective, wider-range wireless communications increases exposure to both passive and active adversaries.

tity management frameworks to IMD environments. However, they must balance these systems with, for example, resource limitations and the size of the trusted computing base.

Accountability

Although preventing malicious activities at all times is impossible, it might be possible to deter such activities by correlating them with an external programmer or entity. Specifically, all IMD setting modifications, as well as all data accesses, could be recorded securely on the IMD (in addition to any logs stored on the programmer)—that is, in a cryptographic audit log that can't be undetectably modified.[10] Physicians could review this log during clinic visits or after detecting certain anomalies in a patient's care. Although current IMDs keep audit logs for the purposes of investigating potential malfunctions, we haven't seen any public discussion of their cryptographic security.

Such an audit log would, however, be meaningless if an external device could claim any identity it chooses—the serial number of an external device or programmer shouldn't by itself act as an identity. Rather, we recommend that legitimate external devices use secure hardware to store credentials—including private keys and the corresponding signed certificates—and authenticate themselves to the IMDs before each transaction. Together, the secure audit logs and the secure identifiers could let an IMD auditor associate individual transactions with the devices performing them. With second-factor authentication, as we proposed earlier, an auditor could also correlate transactional history with a particular healthcare professional.

To address potential DoS attacks against the audit logs' memory size, the IMDs could periodically offload verifiable portions of the audit log to trusted external devices.

Patient awareness via secondary channels

Some IMDs provide an audible alert to signal battery depletion. We recommend using secondary channels to also inform patients about their IMDs' security status. An IMD could issue a notification whenever it establishes a wireless connection with an external programmer or whenever a critical setting changes. As with secure audit logs, these notifications—beeps or vibrations, for instance—won't directly prevent accidental programming or attacks. However, they might help mitigate accidents or deter attacks because the alerts would inform the patient (and possibly bystanders) of the situation and let them react. Other examples of possible secondary channels include an at-home monitor, watch, or phone—all of which could relay further visual, auditory, or tactile information about anomalous security events. Tensions do, however, remain between the use of these secondary channels and patient privacy.

> It might be possible to deter malicious activities by correlating them with an external programmer or entity.

Authorization via secondary channels

Environmental and other secondary elements could serve as factors in authorization. Many existing ICDs, for example, use near-field communication (such as a wand near the chest) for initial activation. After activation, the physician can program the device from a greater distance for a longer period of time. A programming session's extended range and longevity increase exposure for patients because their IMDs might still be receptive to long-range wireless communications after they leave the clinic. Periodically requiring a resumption of near-field communications between the IMD and an authorized external device might therefore be appropriate.

A second approach is to use the built-in accelerometers already in some IMDs. For example, an IMD could cease wireless communications when it detects that its environment has changed significantly, perhaps because the patient stood up from the examining table or otherwise left the clinical setting. By themselves, both approaches will only limit—not prevent—prolonged exposure to adversarial actions.

A separate approach might be to encrypt the communications between the programmer and the IMD, using an encryption key imprinted on a card or a medical-alert bracelet. Here, visual access to the card or bracelet acts as a secondary authorization channel. This approach might, however, lead to safety concerns if a patient forgets his or her card or bracelet and needs urgent emergency care.

Shift computation to external devices

An adversary might use cryptographic mechanisms to mount a DoS attack against the IMD's processor, communications, or battery. A wealth of research exists on improving network security protocols' efficiency under resource constraints—such as

the AUTHORS

Daniel Halperin is a PhD candidate in computer science and engineering at the University of Washington. His research interests include wireless networking, with a current focus on innovative uses of software-defined radios, and practical security and privacy in the wired and wireless digital and physical domains. He received his BS in computer science and mathematics from Harvey Mudd College. He's a member of the ACM, American Mathematical Society, and Usenix. Contact him at the Univ. of Washington, Dept. of Computer Science and Eng., Box 352350, Seattle, WA 98195; dhalperi@cs.washington.edu.

Thomas S. Heydt-Benjamin is a PhD candidate in computer science at the University of Massachusetts Amherst. His research interests include topics in computer security and privacy as applied to mobile and ubiquitous systems. He received his MS in computer science from the University of Massachusetts Amherst. He's a member of the ACM, the IEEE, the International Financial Cryptography Association (IFCA), and the International Association for Cryptologic Research. Contact him at the Computer Science Bldg., 140 Governors Dr., Amherst, MA 01003; tshb@cs.umass.edu.

Kevin Fu is an assistant professor in computer science at the University of Massachusetts Amherst, where he's the principal investigator of the RFID Consortium on Security and Privacy. His research interests include the security and privacy of pervasive technology—including RFID, implantable medical devices, and file systems. He received his PhD in electrical engineering and computer science from the Massachusetts Institute of Technology. He's a member of Usenix, the IEEE, and the ACM. Contact him at the Computer Science Bldg., 140 Governors Dr., Amherst, MA 01003; kevinfu@cs.umass.edu.

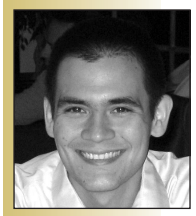

Tadayoshi Kohno is an assistant professor in the University of Washington's Department of Computer Science and Engineering. His research interests include protecting consumer security and privacy in the face of evolving technologies, ranging from electronic voting machines to pervasive healthcare devices. He received his PhD in computer science from the University of California, San Diego. He's a member of the ACM, the International Association for Cryptologic Research, the IEEE, and Usenix. Contact him at the Univ. of Washington, Dept. of Computer Science and Eng., Box 352350, Seattle, WA 98195; yoshi@cs.washington.edu.

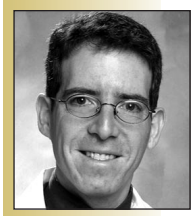

William H. Maisel is the director of the pacemaker and defibrillator service at Beth Israel Deaconess Medical Center and an assistant professor of medicine at Harvard Medical School. His research interests involve the safe and effective use of cardiovascular medical devices—particularly rhythm management devices such as pacemakers and implantable defibrillators. He received his MD from Cornell University Medical College and his MPH from the Harvard School of Public Health. Contact him at the Cardiovascular Division, Beth Israel Deaconess Medical Center, 185 Pilgrim Rd., Baker 4, Boston, MA 02215; wmaisel@bidmc.harvard.edu.

computation offloading via client puzzles[11]—and it's worth exploring how to extend existing DoS limitation methods from conventional networks to IMDs. Although these methods might reduce a DoS attack's efficacy, they might still leave IMDs vulnerable to resource depletion at a lower rate.

Another approach might be to use a resource-rich device to mediate communication between an IMD and an external programmer, much like proposed RFID proxies mediate communication between RFID readers and RFID tags.[12,13] The communication between the IMD and the mediator—perhaps a smart phone, watch, or belt—could use lighter-weight symmetric encryption and authentication schemes, whereas the communication between the mediator and the external programmer could use more expensive asymmetric cryptographic techniques. Increasing the number of devices and protocols involved, however, increases the size of the overall system's trusted computing base, which might make the system harder to secure. For safety, when the trusted mediator isn't present, it might be appropriate for the IMD to fail-open, meaning that caregivers—but also adversaries—could interact with the IMD.

We've proposed research directions for mitigating the tensions between the various goals. However, an ultimate solution will require experts from the medical and security communities, industry, regulatory bodies, patient advocacy groups, and all other relevant communities to collaboratively make decisions on both mechanisms and policies. Our research team is actively exploring the above-mentioned research directions, and we are developing cryptographic and energy-centric methods for providing security and privacy at low cost and without diminishing the efficacy of primary treatments. P

ACKNOWLEDGMENTS

We thank Shane Clark, Benessa Defend, David Eiselen, Barry Karas, Will Morgan, and the anonymous reviewers for their feedback and assistance. US National Science Foundation grants CNS-0435065, CNS-0520729, and CNS-0627529 and a gift from Intel supported this research.

REFERENCES

1. K. Hanna et al., eds., *Innovation and Invention in Medical Devices: Workshop Summary*, US Nat'l Academy of Sciences, 2001.

2. 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," *Proc. Int'l Conf. Parallel Processing (ICPP) Workshops*, IEEE CS Press, 2003, pp. 432–439.

3. US Food and Drug Administration, "Unique Device Identification; Request for Comments," *Federal Register*, vol. 71, no. 155, 2006, pp. 46,233–46,236; www.fda.gov/OHRMS/DOCKETS/98fr/06-6870.htm.

4. T. Drew and M. Gini, "Implantable Medical Devices as Agents and Part of Multiagent Systems," *Proc. 5th Int'l Joint Conf. Autonomous Agents and Multiagent Systems* (AAMAS 06), ACM Press, 2006, pp. 1534–1541.

5. S. Gupta, T. Mukherjee, and K. Venkatasubramanian, "Criticality Aware Access Control Model for Pervasive Applications," *Proc. 4th Ann. IEEE Int'l Conf. Pervasive Computing and Comm.* (PERCOM 06), IEEE CS Press, 2006, pp. 251–257.

6. A. Juels, "RFID Security and Privacy: A Research Survey," *IEEE J. Selected Areas in Comm.*, Feb. 2006, pp. 381–394.

7. M. Jakobsson and S. Wetzel, "Security Weaknesses in Bluetooth," *Progress in Cryptology—CT-RSA 2001: The Cryptographers' Track at RSA Conf. 2001*, LNCS 2020, Springer, 2001, pp. 176–191.

8. M. Gruteser and D. Grunwald, "A Methodological Assessment of Location Privacy Risks in Wireless Hotspot Networks," *First Int'l Conf. Security in Pervasive Computing*, Springer, 2003, pp. 10–24.

9. L. Bauer et al., "Device-Enabled Authorization in the Grey System," *Proc. 8th Int'l Conf. Information Security* (ISC 05), Springer, 2005, pp. 431–445.

10. B. Schneier and J. Kelsey, "Cryptographic Support for Secure Logs on Untrusted Machines," *Proc. Usenix Security Symp.*, Usenix Press, 1998, pp. 53–62.

11. A. Juels and J. Brainard, "Client Puzzles: A Cryptographic Countermeasure against Connection Depletion Attacks," *Proc. Network and System Security Symp.* (NDSS 99), Internet Soc., 1999, pp. 151–165.

12. A. Juels, P. Syverson, and D. Bailey, "High-Power Proxies for Enhancing RFID Privacy and Utility," *Privacy Enhancing Technologies*, LNCS 3856, Springer, 2005, pp. 210–226.

13. M. Rieback, B. Crispo, and A. Tanenbaum, "RFID Guardian: A Battery-Powered Mobile Device for RFID Privacy Management," *Proc. 10th Australasian Conf. Information Security and Privacy* (ACISP 05), LNCS 3574, Springer, 2005, pp. 184–194.