Improving Security of Wireless Communication in Medical Devices

Favyen Bastani and Tiffany Tang

Massachusetts Institute of Technology

Executive Summary

Advancing medical technology has enabled the production of *medical devices* that are externally attached to or implanted inside patients. These devices treat a wide range of conditions, including insulin pumps for diabetes mellitus and pacemakers for various heart conditions. Recently, medical devices have adopted wireless technology to facilitate communication with programmer devices that issue commands to adjust treatment or retrieve sensor data.

While wireless communication has made interaction with medical devices both easier and safer for doctors (for implanted devices, needles previously had to be inserted into patients to carry signals), it has also introduced new security risks. A majority of medical devices implement little to no command authorization and encryption schemes, meaning that malicious attackers can remotely extract sensitive health information from medical devices, or even take control of the device to issue possibly fatal commands.

Security researchers have designed numerous schemes to enforce the security of medical device wireless communication, ranging from shared key derivation algorithms to external wearable devices that handle encryption and authorization. Virtually all such schemes require protocol standardization to make widespread adoption possible, but no standardization organization exists.

Meanwhile, existing regulation gives medical device manufacturers little incentive to improve the security of their devices. Most applicable policy takes the form of guidelines that manufacturers are not required to follow. Since the economic environment pushes for more features and increased device reliability, manufacturers do not adequately consider security when designing devices.

In this paper, we argue that economic, technological, and regulatory factors exacerbate the state of wireless security in current medical devices. From our analysis, we lay out our two-fold recommendations to improving security: first, the Food and Drug Administration (FDA) should require detailed security incident reporting and open investigations following incidents; second, the FDA should encourage the establishment of a multi-stakeholder group tasked with facilitating communication among manufacturers, researchers, health organizations, and regulatory agencies, and standardizing secure protocols.

Table of Contents

Introduction

In today's technologically advanced world, medical devices have become ubiquitous. These devices, including automated insulin pumps and pacemakers that allow many people to lead normal and healthy lifestyles, have grown increasingly complex in design. These devices typically have strong reliability guarantees and are able to operate for years if not decades while implanted in a patient; however, security features such as encryption schemes and user authentication are often not well planned. Even when security features are implemented, careful design choices must be made in order to allow doctors to easily access their patients' devices while keeping potential attackers out. As devices grow more and more intrusive, it is imperative that we focus on providing security guarantees for the individuals that use them.

In fact, security vulnerabilities in wireless-connected medical devices are widespread and severe. Not only do security vulnerabilities threaten patient confidentiality, they also often involve the processing of unauthorized and potentially fatal commands. For example, several pacemakers implement little to no security in wireless communication despite allowing control commands to be transmitted wirelessly, meaning a malicious attacker can easily spoof a command that triggers the vulnerable pacemaker to send deadly shocks to a patient [1]. Attackers can similarly command insulin pumps to manipulate dosage and other settings without the patient's knowledge [2].

Given the vast number of these critical vulnerabilities, the security of wireless communication in current medical devices is clearly inadequate. In this paper, we investigate economic, regulatory, and technological factors contributing to the prevalence of security flaws. Based on this analysis, we then propose policy changes to encourage the adoption and standardization of innovative security defense mechanisms, and to accelerate manufacturers' responses to security threats by requiring detailed security incident reporting to the U.S. Food and Drug Administration (FDA).

One key reason for the continued existence of security issues stems in part from the inherent complexity of achieving security in medical devices. Implanted medical devices (IMDs) are particularly difficult to access physically. When a security vulnerability is disclosed to the manufacturer, applying updates would require extracting the device from the patient's body if the device does not have wireless update capabilities (which is often difficult to implement in small devices). Besides issues with updates, it is also difficult to pack enough computational resources inside implanted and other medical devices to be able to handle the full range of cryptographic operations needed to securely authenticate commands [3].

Even when security vulnerabilities are found, manufacturers often leave them unpatched; instead, manufacturers focus their resources on developing new devices, and may not even maintain development teams for old devices. Increased government regulation is needed to outweigh the economic incentive for manufacturers to ignore security threats; while publicity is generally a sufficient consideration for high-profile industries such as cloud application providers, negative publicity often does not significantly impact medical device manufacturers' sales as patients generally accept whatever devices their doctors prescribe.

Additionally, a large communication barrier between researchers and manufacturers exists, preventing widespread adoption of medical device wireless security innovations. We consider three recent innovations in particular: shielding devices via a wireless jamming system, simplifying encryption of communication between devices by using biometric fingerprints as keys, and offloading computationally intensive cryptographic tasks to an external device. Without coordination between research teams and standardization of inter-device communication protocols, these security systems are not practical.

In Section 1, we assess current regulation of medical devices (primarily by the FDA) and analyze reasons for its ineffectiveness. In Section 2, we consider economic incentives that push companies towards releasing new medical devices with improved reliability and other features, but without incorporating security into their design. In Section 3, we examine the three innovations in medical device wireless security mentioned above and identify why they have not been adopted by manufacturers. In Section 4, we lay out our policy recommendations: first, the FDA should accept security incident reports from both manufacturers and third parties (such as security researchers), and open an investigation for each valid report; second, the FDA should support the establishment of a multi-stakeholder organization tasked with standardizing secure medical device wireless communication systems and facilitating joint efforts among various manufacturers and security researchers. Finally, we draw overall conclusions in Section 5.

1. Current Regulation

Current government regulation of medical devices in the United States primarily consists of the Health Insurance Portability and Accountability Act (HIPAA) and policies enforced by the U.S. Food and Drug Administration (FDA). As we will show, neither of these is adequate to stimulate companies to invest more resources in improving medical device wireless security: HIPAA does not apply to most security vulnerabilities, and FDA regulation is either optional or weak.

1.1 HIPAA

HIPAA was signed into law in 1996 with two distinct goals: to enforce confidentiality of patient health information, especially when transmitted between health care organizations ("accountability"), and to reform health insurance to promote availability of coverage and prevent discrimination based on medical history ("portability") [4].

We will focus on the accountability portion of HIPAA, composed of the Privacy Rule and the Security Rule, which stemmed largely from growing concerns in the 1990's over a possible collision between advancing technology and patient privacy. Note that HIPAA legislation simply required the U.S. Department of Health and Human Services (HHS) to develop regulations in privacy and security areas; thus, in actuality, the Privacy Rule and Security Rule were established by HHS, although today they are commonly referred to as a part of HIPAA.

HIPAA's Privacy Rule set a balance between adequately protecting private health information and ensuring such information can still be accessed for providing quality health care, as well as for public health purposes. It defines *protected health information* as "individually identifiable health information," including health conditions, details of health care, and identifying information recorded along with health records [5]. The Privacy Rule then lays out specific rules defining ways that protected health information may be used, along with mandating that medical record holders provide patients with access to their data [6].

Complementing the Privacy Rule, HIPAA's Security Rule lists standards forming a minimum appropriate level of security safeguards that health organizations must follow when electronically storing or transmitting protected health information [6]. In addition to mandating that organizations conduct ongoing risk analysis to address potential threats to data confidentiality, the Security Rule includes administrative, physical, and technical safeguards that must be followed.

- **Administrative safeguards**: health organizations must establish administrative policies that identify the conditions under which access to data is authorized, and conduct workforce training to ensure compliance with those policies.
- **Physical safeguards**: access to facilities and workstations where protected health information is stored or accessed should be restricted. Records should be kept to track movement of electronic media containing protected information.
- **Technical safeguards**: require implementation of technology to audit accesses to health information, ensure data integrity, and secure data during transmission.

Additionally, HIPAA was extended in 2013 with the omnibus final rule, clarifying that medical device manufacturers are considered business associates and thus also must comply with HIPAA [7].

While HIPAA does succeed in defining how and when personal health information can be stored or transmitted, two critical issues severely reduce its usefulness in regulating wireless communication in medical devices: HIPAA's inapplicability to security issues that don't impact privacy, and a lack of enforcement by HHS.

First, medical device vulnerabilities relating to the acceptance of unauthorized commands, which can lead to patient injury or even death, do not involve exposure of protected data, and thus fall entirely outside the scope of HIPAA [8]. Thus, even though HIPAA's Security Rule specifies cryptographic techniques that should be used in network transmissions, manufacturers producing medical devices that do not transmit sensitive data are not currently specifically required to utilize these techniques. Furthermore, since HIPAA's primary goal is to curb privacy issues relating to health organizations' use of personal medical records, an amendment that adds unrelated regulation over security vulnerabilities distracts from that goal.

Second, HIPAA enforcement actions are rare. The first enforcement action by HHS that involved concrete penalties did not occur until 2008 (over a decade after HIPAA was passed), when HHS found that Providence Health & Services employees frequently brought devices storing confidential data home, and in some cases these devices were stolen. HHS and Providence reached a settlement where Providence paid a fine of $100,000 and was required to follow a corrective action plan to resolve issues in its privacy policies [9]. Although since then HHS has yielded several more settlements, these have primarily targeted large health organizations and enforcement gaps still remain [10].

1.2 Food and Drug Administration

FDA regulation of medical devices primarily consists of a premarket approval application process, guided by the Medical Device Regulation Act (which in 1976 amended the Food, Drug, and Cosmetics Act). Medical devices are grouped into three classes based on risk, with Class I consisting of low-risk devices such as tongue blades and Class III consisting of high-risk devices such as insulin pumps and pacemakers [11]. We focus on Class III devices as they pose the largest danger to patients from a wireless communication vulnerability standpoint.

Manufacturers have two options to receive FDA approval for Class III medical devices. For devices "substantially equivalent" to already approved devices, the manufacturer can simply show that the new medical device approximates the existing device in both usage and technical specifications. Otherwise, the manufacturer must undergo an analysis process to demonstrate the safety of the medical device; in this case, the FDA may require data such as results from clinical trials and laboratory testing [12].

Several issues with this regulatory architecture limit its effectiveness in ensuring wireless communication security. First, in the premarket evaluation process, the FDA focuses on evaluating the medical device's ability to operate correctly under various conditions and extended timeframes. While this is important to avoid patient injury from the medical device itself, it ignores security problems that could have even more severe consequences when a remote attacker is able to exploit a vulnerability on a large scale. One reason for this gap in the FDA evaluation is that the FDA hires very few software engineers who would be knowledgeable of design practices that promote security [13].

Additionally, many medical devices unreasonably pass the "substantially equivalent" test. For example, devices that replace hardware systems with software often still are approved without an analysis process [13] despite incidents like Therac-25 [14] that show that such replacements frequently introduce critical bugs. Similarly, the addition of a wireless communication component exposes devices to remote exploitation but frequently does not lead to an extended approval review. Strikingly, a 2011 study found that over 70% of high-risk medical device recalls issued by the FDA between 2005 and 2009 involved devices that were originally approved as "substantially equivalent" to a previously approved device [12].

The FDA has recently made progress to remedy the lack of security evaluation of most medical devices. In October 2014, the FDA released guidelines specifying actions that medical device manufacturers should take to ensure maximum security. However, the FDA specifically states that the guidelines "do not establish legally enforceable responsibilities," and therefore largely consists of cybersecurity information that manufacturers are already aware of (for example, that devices should implement strong authentication policies, code signing to authorize updates, and logging of device commands) [15]. As a result, these guidelines will likely have little effect on the current security situation.

2. Economic Incentives

This section presents the economic incentives and motivations of medical device manufacturers, and reasons why current practices have persisted for so long. We show that current incentives are insufficient to encourage manufacturers to invest resources in improving the security of medical devices.

2.1 Trends of the United States Medical Device Industry

The United States medical device industry is one of the most prosperous in the world, with a market size of $110 billion and more than 6,500 companies [29]. Top companies such as Johnson & Johnson, General Electric Co., Medtronic, Inc, and Siemens AG generate billions of dollars in revenue each year [30]. According to a study in 2008 on the economic impact of the medical technology industry by the Lewin Group, the industry employed 422,778 workers, paid $24.6 billion in earnings and shipped $135.9 billion worth in products [33]. This booming industry has led to a large focus on the development of novel medical devices and reliability features, but only at the expense of maintaining security updates for existing devices.

In the current medical device economic ecosystem, manufacturers have a natural tendency to produce and ship devices as quickly as possible. Maintaining the wide range of old devices that a manufacturer previously produced is costly, especially when resources can be diverted to building newer devices that generate concrete profit. So, without an incentive that makes this maintenance expense more desirable than having unpatched security flaws, manufacturers will tend to ignore issues and push for features.

Consequently, the medical device industry has generally failed to address existing security vulnerabilities. Both professional security researchers and part-time hackers frequently discover security flaws in medical devices. For example, for the 2011 BlackHat conference, Jerome Radcliffe published a paper demonstrating an attack where an implanted insulin pump is misled into accepting unauthorized commands [31]. His paper contains detailed instructions and code. Also in 2011, Dina Katabi and her research group at MIT revealed a similar vulnerability targeting pacemakers that can trigger deadly voltages in the patient. Thus far, neither vulnerability has been fixed.

In general, the medical device industry has a slow tendency to adapt to change. We show in the next sub-section why this is the case.

2.2 Preemption Clause

The *doctrine of preemption* reduces legal pressure on manufacturers by protecting manufacturers from being held liable in certain cases where devices are found to be defective (here, this defect would be a security vulnerability). The doctrine of preemption, which is under the Supremacy Clause in the U.S. Constitution, states that federal law preempts state law. Thus, the Medical Device Regulation Act mentioned in Section 1 takes precedence over any state law.

In a paper which analyzes preemption and liability of medical devices, the authors state that medical device manufacturers in general are not held liable for medical device safety [28]. Thus, if even medical device safety is not a primary concern, then medical device security most definitely will not be a consideration. And unfortunately, the "majority of State courts have shown strong support for *Riegal* [*v Medtronic, Inc*, the case that enforced the Medical Device Regulation Act] and its progeny, and the proposed legislation to reverse the Supreme Court decision has gathered a relatively small fraction of co-sponsors in Congress [28]."

There are two important details to note. First, as discussed in Section 1.2, if a manufacturer wants to change the design of their medical device such that it affects safety and usability the device must gain premarket approval again. This might discourage manufacturers from going through the process of gaining approval again. Secondly, as pointed out in Section 2.1, the medical device industry generates a lot of revenue for the United States. Stricter regulations by the government on the industry may harm its growth, which might be a concern for lawmakers.

A prominent example of this as well as the tendency for government to support the preemption clause is shown in the case *Stengel v Medtronic*. The incident occurred in 2005, in which a patient's lower half was was paralyzed due to the use of a pain pump manufactured by Medtronic. The patient was not warned about the possible complications for using the device. The Ninth Circuit Court initially ruled that the Medical Device Regulation Act preempt the state safety laws the plaintiff claimed that the device violated. However, on January 10, 2013, the concept of "parallel claims" was introduced, which stated that the Medical Device Regulation Act does not preempt state-law failure to warn claims [34]. Medtronic was denied a writ of certiorari on June 2014 (which also shows that the medical device company is not particularly interested in patient safety) [35]. However, the Court has left the concept of "parallel claims" to be debated as it has not been clearly defined.

2.3 Design Process

As a result, the design process of hardware and software is poorly thought out in these devices. According to Bruce Schneier, a security technologist and author of several essays on security and technology, there is little consideration on maintaining or supporting devices after they are released out in the market during the device design process [16]. Not surprisingly, the main focus of most companies is on profit. The cheapest chips are the ones that will be used in medical devices that people will have to depend on. Software on medical devices is often outdated by a few years. Patches and updates are generally difficult to accomplish, if not impossible, because the device contains only the binary and not modifiable source code; without such source code, updates to the device take significantly more time to write.

Aside from the lack of concern on security in the medical device industry, there are other reasons why this problem persists. During the FTC's Internet of Things workshop, Craig Heffner, a vulnerability researcher, claims that neither vendors nor consumers care about security [17]. If the users themselves do not consider security seriously, then there will be even less incentive for manufacturers to incorporate security when designing devices.

Finally, we take into consideration the current regulation of security of medical devices. As stated earlier in this section, companies are less likely to be held liable in the case that a customer falls victim to a security breach. Without the needed pressure from government regulation, businesses do not have to worry about the repercussions brought on by not tightening security. In October 2014, the FDA published a list of non-binding recommendations on medical security. Later in the month, US Homeland and Security investigated the security of medical devices and discovered several security-related issues. Hospira Inc, Medtronic Inc, and St Jude Medical Inc "said they take cybersecurity seriously and have made changes to improve product security, but declined to give details [32]." Even though the security flaws are publicized, there is no assurance from the medical device industry that security will be improved, which strongly supports the idea that there is not enough pressure from both the government and the public.

In Section 4.1, we present a method to fix reported security vulnerabilities within a reasonable time period. This solution will be regulated by the FDA and does not conflict with either the preemption doctrine or the Medical Device Regulation Act. It does, however, incentive medical device companies to focus on security by exerting appropriate government pressure.

3. Wireless Security Technology

Security and network communication researchers have developed various innovative techniques to improve the security of wireless communication in medical devices. Recent work has yielded schemes that derive encryption keys from biological signals to easily but securely synchronize multiple devices on the same patient, and external devices that protect implanted medical devices (IMDs) from unauthorized commands in various ways. However, as of yet, none of these innovations have achieved mainstream implementation. We explore three specific innovations in more detail, and find that all depend on some degree of standardization of medical device wireless security systems so that devices from different manufacturers can interoperate; however, no organization exists to facilitate this standardization. Additionally, there are communication gaps between various stakeholders, particularly between security researchers and device manufacturers, which result in security innovations that do not entirely align with manufacturers' needs.

3.1 Keys from Biological Fingerprints

In a 2006 study, Poon et al. tackled the problem of body area sensor networks (BASN) where various medical devices, some implanted and some external, need to wirelessly communicate with each other to improve treatment [18]. For example, an insulin level monitoring device could report to an insulin pump when it detects low insulin levels.

Because of the interconnected nature of a BASN, where devices can influence each other either by providing sensor data or by directly sending control commands, encrypting wireless communication is critical to prevent the unauthorized injection of malicious data. The researchers point out that encryption schemes also mitigate privacy concerns of health data transmitted between devices, and ensure that devices of individuals in close proximity do not interfere with each other.

However, typical cryptographic strategies are difficult to apply to BASN. First, key agreement protocols such as Diffie-Hellman and RSA are computationally expensive and thus often impossible to implement on small medical devices [19]. Still, a shared key between devices that need to communicate must somehow be produced in order to use less expensive symmetric encryption algorithms. While this key could be hardcoded into the medical devices, generally medical devices need to be modular as each patient may be prescribed a different set of devices, and patients may also add new wearable devices that need to know the same key; a hardcoded key scheme also itself presents security issues. Furthermore, hardcoding the key leads to communication difficulties when interacting with programmer devices that doctors may use to tune a patient's medical device configuration; moreover, in emergency situations, it may be infeasible to have the programmer device be preconfigured with the key used by the patient's medical devices.

Other approaches have been proposed to solve parts of the problem, but generally either require expensive asymmetric encryption or maintain the fixed shared key requirement. For example, Rasmussen et al. developed an access control scheme based on confirming proximity, where ultrasonic distance bounding protocols are used to ensure external devices are within some threshold in physical distance to an IMD [20]. To handle the case of emergencies, the researchers allow programmer devices that don't have knowledge of the key to still send commands if they demonstrate via the proximity protocol that they are within a few centimeters of the patient. However, in their scheme, a shared key still needs to be propagated across the medical devices, and devices need to be outfitted with additional ultrasonic circuitry.

Poon et al. propose to derive keys from biological fingerprints instead of having them be preset [18]. This form of key derivation, which ensures that each of one patient's medical devices are able to derive a single shared key, facilitates symmetric encryption between the patient's devices without the potential to leak keys to attackers who do not have direct physical access to the patient.

The researchers consider various signals including glucose levels, heartbeat timing, and blood pressure. They find that heartbeat timing information produces sufficient entropy to prevent attackers from easily predicting the random bits, and at the same time is a signal that is readily available to most medical devices. Thus, they develop a scheme where the interpulse interval is measured with electrocardiogram or photoplethysmogram data, and then use those measurements to establish a key.

A wireless time synchronization scheme must first be followed to ensure that all devices begin monitoring interpulse interval at the same time; each device then monitors for a predetermined duration and at a specific sampling frequency, e.g. 200 Hz. Once the sampling process completes, a binary encoder merges the samples into a 128-bit binary sequence, which is used to decide on the symmetric encryption key.

Due to noise, different medical devices might not arrive at exactly the same binary sequence. To handle this, devices can use a fuzzy commitment scheme [21], which provides an algorithm that accepts a random codeword and a key derived from biometric data and outputs a commitment that encodes the codeword under the key in such a way that the original data cannot be extracted without the key. Since the commitment scheme is "fuzzy", decommitting the codeword can be done as long as the Hamming distance between the actual key and the corresponding key determined at another device differs by less than some threshold number of bits. Using this scheme, two medical devices can establish an encrypted communication channel, through which sensor data and control commands can be sent securely, by using the codeword as a symmetric encryption key. Although the commitment scheme allows some variability in the key, the channel is still secure since an attacker attempting to use values with Hamming distance outside the fuzziness threshold cannot extract any information about the codeword.

There are, though, practical limitations to the key derivation scheme that prevent its adoption by medical device manufacturers. First, while the interpulse interval provides enough entropy after a suitable duration, this duration may in fact be on the order of minutes. In many situations, it is imaginable that a medical team will need to immediately send emergency commands to the devices in a patient, and in those cases waiting minutes for the key agreement process to complete would be unacceptable. Also, due to exponentially rising computational power of standard hardware, the amount of time needed to gather enough entropy from interpulse interval to produce secure keys may increase as attackers are able to more quickly guess encryption keys.

Additionally, while most medical devices can readily measure heartbeat timing, some kinds of devices may not have access to that information, and extending those devices to conduct their function while also measuring heartbeat timing may greatly increase their complexity and cost.

Studies have also demonstrated that surface-penetrating radars (also used in detecting unexploded mines or structural defects in buildings) can measure human heartbeat data [22]. While these techniques so far do not have the precision to conduct measurements on the level of accuracy needed to steal the key (which is based on data sampled at 200 Hz and with a very precisely coordinated sampling initiation time), future advancements in this area are likely.

Moreover, novel key agreement protocols like this that are specifically tailored to use in medical devices need standardization so that devices from different manufacturers can still communicate. Such standardization organizations simply do not exist.

3.2 IMDShield

Rather than attempt to enable medical devices to more easily implement cryptographic techniques, IMDShield [23] aims to use wireless signal jamming to protect IMDs from unauthorized eavesdropping and commands. Programmer devices are only allowed to communicate with the IMD via a shield device that handles any necessary authorization and encryption with the programmer device. The shield device then forwards the command to the IMD. The shield device begins jamming when it detects an unauthorized signal to the IMD, or when the IMD responds to a query (such as a programmer device requesting the IMD to upload sensor records).

The shield device has the ability to simultaneously transmit and receive wireless signals. It continuously checks for signals that match the header of an IMD command packet (or some other sequence denoting a command that IMDs may pick up), which represent unauthorized signals since legitimate programmer devices should always communicate with the shield device on a separate frequency and with encryption. When it finds a matching signal, the shield device immediately begins jamming communications so that the actual command packet content does not reach the IMD.

When the shield device does receive an authorized command, it will forward the command to the IMD when jamming is not enabled (attackers could still mount denial of service type attacks by ensuring the shield device has to constantly emit the jamming signal); if an unauthorized command is detected while the shield device is transmitting the authorized command, the shield will have to cancel the transmission and switch to jamming to ensure no malicious actions get through. After the shield device finishes transmitting, it will start jamming to ensure the confidentiality of the response from the IMD; the jamming should stop after a duration corresponding to the longest possible IMD response. Note that this means that the shield device must be able to receive data from the IMD even while it is jamming.

In order to receive while jamming is in progress, IMDShield employs a unique system where the shield device emits both a random jamming signal and an antidote signal; the antidote cancels the jamming signal at the receiver antenna of the shield device only; at any other location, the jamming signal combines with the IMD response signal, making it impossible to extract the response. This allows the shield device but no other receiver to retrieve ungarbled data from the IMD.

IMDShield is promising in that, unlike other solutions to attaining wireless communication security, it does not require modification to IMDs. However, it still presents usage and security issues under some conditions.

First, using an extra device to manage security alleviates the complexity of implementing secure encryption techniques on very constrained medical devices, but at the same time burdens users. Patients must remember to wear the shield device at all times, or they will be vulnerable to malicious signals. The shield device must also be periodically recharged. Additionally, programmer devices will need to support an additional protocol to communicate with the shield device: if they only support communicating through the shield device, problems could emerge in emergencies where the shield device was removed (for example, it may need to be removed if jamming activates incorrectly). While programmer devices are not as constrained as IMDs, this does nevertheless raise manufacturing complexity.

Furthermore, the use of a jamming signal poses a wide range of concerns. For example, when patients with similar medical devices are in close proximity and a doctor is attempting to program one set of devices, the jamming system of the other patient might activate and prevent programmer operation. Similarly, if jamming activates outside the isolated environment of a hospital (which may happen by chance depending on the length of the IMD packet header, since noise could be misinterpreted by the shield device as a malicious signal), it may interfere with other normal activities; even in the hospital, other systems could be affected. Also, note that under IMDShield, communication between one patient's IMDs is not possible since the shield device cannot distinguish between malicious traffic and traffic originating from one of the patient's own IMDs.

The need for the shield device to have knowledge about the IMDs is also problematic. At a minimum, the shield device needs a header sequence to know when to activate the jamming signal against malicious commands, and a maximum IMD response time to know when jamming of IMD responses should be disabled. On the other hand, patients may have multiple implanted or wearable medical devices that all need to be protected, complicating the shield device design. This also mandates manual input of the technical details when a new medical device needs to be registered, which could lead to human errors; the process could be automated if medical devices broadcast their specifications, although this presents its own security concerns (for example, an attacker could register a fake device with a long response time and with a short and common header, to cause the shield device to frequently activate jamming).

Lastly, since IMDShield is based entirely around jamming, malicious signals with enough power can still bypass the shield and act on the IMD (specifically, the attacker would need one hundred times the jamming power of the shield device).

3.3 IMDGuard

IMDGuard [24], designed by Fengyuan Xu et al., combines ideas from the heart-rate-based key derivation system in 3.1 and IMDShield in 3.2. Complex and computationally intensive cryptographic tasks needed to decrypt and authorize incoming commands from remote programmer devices are offloaded from IMDs to an external guardian device. Thus, similar to IMDShield, IMDGuard acts as an authentication service between programmer devices and IMDs. Rather than relying exclusively on jamming to block unauthorized transmissions, though, the IMDGuard system encrypts communication between the guardian device and IMDs using a shared key derived from electrocardiogram signals; in some circumstances that we will discuss, IMDGuard does still require low-power jamming.

In the IMDGuard scheme, the programmer begins by sending a session initialization request to the IMD. The IMD generates a nonce (a randomly generated number used to prevent replay attacks) and sends the nonce with its unique identifier to both the programmer and the guardian. After the guardian receives this message, it sends a nonce to the programmer, and the programmer sends back the same nonce signed with the programmer's private key. The guardian, which should already be configured with the programmer's public key, verifies this signature.

If the guardian successfully authenticates the programmer, the guardian generates a shared session key for communication between the programmer and the IMD. This is securely relayed to the IMD by encrypting with the shared key (and including the nonce value), and to the programmer by encrypting with its public key. If authentication fails, the guardian will notify the IMD to reject the transmission.

The shared key is derived with a system similar to that described in Section 3.1. However, instead of using fuzzy commitment or a similar scheme, IMDGuard quantizes the interpulse interval data from electrocardiogram samples in such a way that there is a low probability of mismatch.

In emergency situations, the IMD may have to be programmed without going through the authorization protocol (the programmer being used may be one that the guardian is not aware of). In these cases, we can assume that the patient is in a secure facility and there won't be unauthorized transmissions. Thus, we want medical personnel to be able to remove the guardian device and then program the IMD without encryption.

In IMDGuard, this goal is achieved via an emergency condition protocol where the IMD confirms that the guardian is not present. First, when waiting for an authorization success or failure message from the guardian after the programmer makes a session request, the IMD times out after some maximum waiting duration. However, the IMD not receiving a message from the guardian does not imply the guardian is absent as an attacker may be jamming the guardian's signal. So, the IMD does proceed with communicating with the programmer, but the guardian device will jam these messages if the guardian is still present.

Specifically, the IMD sends an initial challenge nonce to the programmer after activating the emergency condition protocol (following timeout while waiting for guardian authorization response), waits some time, and then sends another nonce. If the programmer receives both nonces, it computes a predetermined function on the two nonces and returns the result to the IMD. The IMD validates the result, and if successful, accepts an unencrypted command from the programmer (and provides an unencrypted response when applicable). To prevent attacks where the programmer jams the guardian, if the guardian is present and sees the initial nonce of the emergency protocol being sent, then the guardian will jam the second nonce so that the programmer cannot compute a valid result.

While IMDGuard still depends on jamming to some extent for security, this jamming blocks the IMD's response rather than the attacker's signal. Since the attacker cannot control the strength of the IMD's communication, this jamming can be low-power and still remain secure against any attacker. The scheme also supports communication between multiple IMDs on a single patient, with the heart-rate-based shared key.

IMDGuard does require that each IMD be designed to interface with the guardian device, meaning modification to the IMD is necessary. Also, the shared key is vulnerable to the same attacks discussed in 3.1.

3.4 Standardizing Protocols and Facilitating Collaboration

Both the biological signal key agreement system in Section 3.1 and the IMDGuard device in Section 3.3, besides having some security flaws that may be resolvable with technology, require the establishment of standardized protocols to facilitate secure communication between the multiple medical devices. Without such a protocol, only devices produced by the same manufacturer would be usable in conjunction, posing problems of vendor lock-in and incompatibility with old devices. This is particularly concerning for IMDGuard, since requiring a unique type of external device to protect each IMD makes the solution much less feasible; patients with multiple IMDs operating independently would need to continuously wear multiple shield devices, and health organizations would have to decide quantities of different shield device brands to purchase.

While IMDShield does not require modifications to IMDs, a single shield device should still be able to support management of multiple IMDs for the same reasons as above. Standard protocols would be needed to achieve this, as the shield device needs to know the IMD command headers to cross-check for in incoming signals to determine when the shield should start jamming. Additionally, the shield device requires a standard communication protocol to interact with various programmer devices.

Another issue common to all of the innovations discussed above is that there has been little movement in adoption by manufacturers. Most research projects in this area are carried out without significant communication with manufacturers, resulting in technologies that manufacturers are not willing to implement—for example, the dependence on jamming in IMDGuard is particularly concerning.

In Section 4.2, we will argue that a multi-stakeholder organization consisting of both medical device manufacturers and security researchers would be needed to conduct the standardization process. Besides presenting an opportunity to develop common protocols that can be used across medical devices, such an organization also facilitates communication not only of security innovations to manufacturers, but also between research groups. This may yield additional research in areas closer to what manufacturers are willing to produce; for example, it may be possible to integrate the functionality of an external device like IMDShield or IMDGuard with a device that has existing medical applications so that an additional separate device, which could burden patients, would not be needed.

4. Recommendations

Our recommendations are two-fold. First, to ensure that medical device manufacturers take security threats seriously, we propose an approach modeled on the National Highway Traffic Safety Administration’s procedure for car recalls; specifically, we recommend increasing reporting requirements for security threads and adopting regulations that apply penalties when manufacturers do not adequately respond to these threats. Second, the FDA should encourage the establishment of a multi-stakeholder organization that includes security researchers, medical device manufacturers, health organizations, and regulatory agencies tasked with developing promising medical device wireless security systems and standardizing the systems so that devices may interoperate.

4.1 Responding to Security Vulnerability Disclosures

We propose a response system that would require medical device manufacturers to resolve vulnerabilities within a certain timeframe.

It is inevitable that software will contain bugs, so we do not expect manufacturers to produce flawless code. However, manufacturers do have a responsibility to patch security flaws that pose a high safety risk to patients, including any vulnerabilities that can be exploited remotely to cause patient injury. Therefore, in the case of a security breach, manufacturers should be required to fix the bug and provide updates to their customers until the vulnerability no longer poses a threat. The following process illustrates a possible regulatory framework that the government can adopt to encourage medical device manufacturers to respond to security vulnerabilities in a timely fashion. We base our solution on the National Highway Traffic Safety Administration’s procedure for car recalls. We argue that the FDA should administer investigations because the FDA already has the power to conduct medical device recalls. In addition to security issues, our solution would enable the FDA to act on other defects such as high-risk usability issues.

We begin by defining the security goals. *Medical device security* is the ability of a medical device to protect the user against safety threats that may "pose unreasonable risk of death and injury" [25]. These *safety threats* include security vulnerabilities that can be trivially exploited, along with other design flaws that threaten a patient's health.

In general, the FDA should be informed about any discovered safety threats. In the case that the medical device manufacturer first discovers the safety threat, the manufacturer should update the software to remove the safety threat. If an entity not affiliated with the manufacturer discovers the threat, such as a security researcher or a consumer, then that entity should file an incident report to the FDA.

After the FDA receives an incident report, the FDA will open an investigation, involving a careful examination of the safety safety threat. In the case of security vulnerabilities, the FDA would consult third party security specialists. We outline the investigation procedure below.

Initial Screening: During this phase, the FDA examines incident reports. If there is enough reason to believe that there is a vulnerability that threatens medical device security, then an investigation will be opened.

Petition Analysis*:* In this phase, the FDA formulates a response to the entity who filed the initial incident report. Here, the FDA may choose to open an investigation, or to decide that there is not enough evidence. In the latter case, the FDA will provide reasons why an investigation will not be opened to the incident reporter; depending on the nature of the threat, the incident reporter may be asked to provide more details on the threat.

Investigation*:* The software and hardware of the medical device will be examined by security experts for security vulnerabilities. The medical device manufacturer would be required to provide the FDA information to speed up the investigation process, including the entire software codebase and a detailed description of the hardware. During this investigation, any bugs found by the FDA (or third parties that the FDA consults) must be patched by the medical device manufacturer.

Effectiveness Study: At the end of the investigation, the FDA will draw conclusions on the effectiveness of the entire incident for future investigations.

If the FDA determines that there is a security threat, then the medical device manufacturers are required to issue a patch for the vulnerability within forty-five days. In the case of a hardware vulnerability, the defective medical device will be replaced. The patch or a replacement device must then be provided free of charge to the consumer.

Any individual injured or harmed due to the vulnerability may additionally take independent legal action against the medical device manufacturer.

Some medical device manufacturers will oppose the adoption of such a policy. First, the proposed solution would force manufacturers to respond to security vulnerabilities in a timely fashion, or risk further penalties; as a result, companies will inevitably have to allocate more resources towards hiring security experts in order to be able to understand and fix security issues. However, this outcome in fact makes it critical for the solution to follow through, as these resources are necessary to improve medical device. Second, manufacturers have traditionally kept medical device technologies highly proprietary, and thus would be against any solution that forces disclosure of these technologies, even if only to the FDA. This concern may be mitigated by allowing the manufacturer to select a third party to conduct an external security audit (subject to FDA's approval).

Still, many other groups would be in favor of this policy, including security researchers, who dedicate significant amounts of time towards discovering security flaws, and consumers, who depend on these devices and are at risk until existing security vulnerabilities are resolved. There may in fact be some medical device manufacturers who are in favor of this proposed solution. The added assurance that security vulnerabilities will be resolved within a reasonable time period may encourage consumers to purchase their medical devices from companies that follow this protocol. It grants the medical device manufacturers credibility, which is an important aspect in building a large consumer base.

4.2 Multi-stakeholder Organization

We further recommend that the FDA encourage and support the establishment of an open governance multi-stakeholder organization that both serves as a forum for exchange of ideas between research groups and medical device manufacturers, as well as facilitating standardization of wireless communication protocols to enable security in complex and small devices. Under open governance, any individual would be able to participate in meetings of the organization.

Currently, medical device technology is kept highly proprietary [13]. With closed and undocumented network protocols being the industry norm, security researchers have limited ability to evaluate the effectiveness of security controls. These trends also impede interoperability between medical devices of different manufacturers, and push up costs for health organizations (and thereby costs for patients as well) as different sets of devices from multiple manufacturers need to be purchased.

Similarly, many of the most cutting-edge innovations in medical device wireless security are feasible only when devices share a common communication framework. Offloading security protections to an external device, as suggested by both IMDShield in Section 3.2 and IMDGuard in Section 3.3, cannot be implemented on a wide scale if hospitals must purchase a different external device model to protect each unique medical device: doctors would need to be trained in the quirks of each such protection device and how they should be used in conjunction with the corresponding programmer device, there would be a greater likelihood of interference between different protection devices, and patients with multiple medical devices would need to keep track of wearing and charging several external devices. Similarly, the key agreement schemes described in Section 3.1 require standardized protocols in order to establish inter-device communication.

Furthermore, while medical device manufacturers have excelled in reaching out to doctors to explain what features medical devices need to make them work well, both when doctors need to adjust configuration and when patients need to verify their device's status, there is a gap in communication between manufacturers and security researchers. Indeed, as Kevin Fu points out [13], just as the FDA lacks the security experts to identify vulnerabilities in medical devices, manufacturers may also lack such experts and thus not be aware of fundamental design decisions that threaten security.

Multi-stakeholder organizations have successfully tackled similar situations where non-standard technology would undermine usability, security, or interoperability. Most notably, the Internet Engineering Task Force (IETF) and other groups overseen by the Internet Architecture Board promote standard Internet protocols that enable intercommunication between networks, including Internet Protocol, transport layer protocols, and routing protocols [26]. For example, the Border Gateway Protocol (BGP) enables networks operated by distinct organizations to communicate with each other to achieve efficient and stable routing. Its design involved collaboration between network researchers, operators, and other IETF members that resulted in a protocol capable of scaling to large numbers of sub-networks while still providing optimal routes. Edge router manufacturers specifically design their routers to support protocols like BGP, adding hardware optimizations that allow high performance and maintenance of large routing tables. Without a standard protocol like BGP, establishing interconnection arrangements (such as peering) would be much more expensive as the involved parties would need to identify a routing protocol for the interconnection and purchase new routers for that protocol; with BGP, they would only need to link existing routers.

Likewise, the International Medical Products Anti-Counterfeiting Taskforce (IMPACT) is a multi-stakeholder organization that seeks to "to improve coordination and harmonization across and between countries so that eventually the production, trading and selling of fake medicines will cease" [27]. By combining expertise from pharmaceutical manufacturers, health regulatory agencies, national customs organizations, and other groups, IMPACT has identified key problems enabling counterfeit medicine producers to avoid prosecution and blocking trade of such products, and continues to work with both international and national enforcement authorities.

Security researchers would benefit from a multi-stakeholder organization as it could allow them access to more closely investigate existing wireless security designs in medical devices, and how they can be improved; additionally, it would open opportunities for joint efforts with manufacturers and health organizations to develop new wireless transmission security innovations that fit with the medical device model. Health organizations would be able to offer better health care when patients are not concerned about the security implications of implanted medical devices that allow commands to be sent over wireless signals.

While device manufacturers might resist the adoption of open protocols due to industry traditions, this would be outweighed by the increased security of their products that would improve not only patient safety, but also their reputation. This is especially important as news articles more and more frequently cover discoveries of critical vulnerabilities in medical devices.

Overall, establishing a multi-stakeholder group consisting of security researchers, medical device manufacturers, health organizations, and regulatory agencies, with the goals of specifying standard, open, and secure communication protocols and facilitating greater interactions between all parties, would be a significant step towards improving security. Protocols and systems produced in cooperation between manufacturers and researchers have the potential to not only be secure, but also to meet practical demands of hospitals and patients. Additionally, with the adoption of standard, open protocols, security evaluations of medical device systems can produce much more concrete results, and both security protection devices and programmer devices will become simultaneously more affordable and more secure as competition grows.

5. Conclusions

While medical device manufacturers are making significant progress in improving the reliability of devices in normal operation, the security of wireless communication in these devices has not received as much attention. Yet, vulnerabilities often allow attackers to take full control of devices and perform actions that may gravely injure patients.

To remedy the situation, we propose policies that a) ensure resolution of security threats by requiring detailed reporting of security problems and punishing companies that fail to respond within a reasonable amount of time; and b) facilitate communication between security researchers and the medical device manufacturing industry (along with other stakeholders) by establishing a multistakeholder organization. While these recommendations do not solve the complexity of improving security in medical devices, they do put into place policies that will incentivize development of security techniques.

References

[1] "Pacemaker hack can deliver deadly 830-volt jolt," Jeremy Kirk, ComputerWorld, 17 October 2012.
[2] "Researcher battles insulin pump maker over security flaw," Ellnor Mills, CNET, 26 August 2011.
[3] "Inside Risks: Reducing Risks of Implantable Medical Devices," Kevin Fu, *Communications of the ACM*, 52 (6), June 2009.
[4] "The Politics of the Health Insurance Portability and Accountability Act," Brian K. Atchinson and Daniel M. Fox, *Health Affairs*, 16, 1997.
[5] Health Insurance Portability and Accountability Act of 1996.
[6] "Challenges Associated with Privacy in Health Care Industry: Implementation of HIPAA and the Security Rules," Young B. Choi et al., *Journal of Medical Systems*, 30 (1), June 2005.
[7] "Modifications to the HIPAA Notification Rules," Department of Health and Human Services, 25 January 2013.
[8] "Improving the security and privacy of implantable medical devices," William H. Maisel and Tadayoshi Kohno, *New England Journal of Medicine* (362;13), 1 April 2010.
[9] "HIPAA Security Enforcement is Here," Kirk J. Nahra, *IEEE Security & Privacy*, November 2008.
[10] "HHS Case Examples and Resolution Agreements," U.S. Department of Health & Human Services, <http://www.hhs.gov/ocr/privacy/hipaa/enforcement/examples/> (accessed 6 November 2014).

[11] “Medical Device Regulation: An Introduction for the Practicing Physician,” William H. Maisel, *Medicine and Public Issues*, 2004.
[12] “Medical Device Recalls and the FDA Approval Process,” Diana M. Zuckerman et al., *Arch Intern Med*, 14 February 2011.
[13] “Trustworthy Medical Device Software,” Kevin Fu, *Institute of Medicine Workshop on Public Health Effectiveness of the FDA*, 11 April 2011.
[14] “An Investigation of the Therac-25 Accidents,” Nancy G. Leveson and Clark S. Turner, *IEEE Computer*, July 1993.
[15] “Content of Premarket Submissions for Management of Cybersecurity in Medical Devices,” U.S. Food and Drug Administration, 2 October 2014.
[16] “The Internet of Things is Wildly Insecure--And Often Unpatchable,” Bruce Schneier, Wired, 6 January 2014.
[17] “FTC Internet of Things Workshop,” Federal Trade Commission, 19 November 2013, Transcript.
[18] “A Novel Biometrics Method to Secure Wireless Body Area Sensor Networks for Telemedicine and M-Health,” Carmen C. Y. Poon et al., *IEEE Communications Magazine*, April 2006.
[19] “Security and Privacy for Implantable Medical Devices,” Daniel Halperin et al., *Pervasive Computing*, January 2008.
[20] “Proximity-based Access Control for Implantable Medical Devices,” Kasper B. Rasmussen et al., *ACM Conference on Computer and Communications Security*, November 2009.
[21] “A Fuzzy Committment Scheme,” Ari Juels and Martin Wattenberg, *ACM Conference on Computer and Communications Security*, November 1999.
[22] “Detection of Human Breathing and Heartbeat by Remote Radar,” S.I. Ivashov et al., *Progress in Electromagnetic Research Symposium*, 28 March 2004.
[23] “They Can Hear Your Heartbeats: Non-Invasive Security for Implantable Medical Devices,” Gollakota et al., ACM SIGCOMM 2011, 15 August 2011.
[24] “IMDGuard: Securing Implantable Medical Devices with the External Wearable Guardian,” Fengyuan Xu et al., IEEE INFOCOM 2011, 10 April 2011.
[25] Motor Vehicle Safety Defects and Recalls, NHTSA, May 2011.
[26] “The Internet, its Governance, and the Multi-Stakeholder Model,” Richard Hill, *Info*, 28 November 2013.
2d
[27] IMPACT Brochure, World Health Organization, May 2008, <http://www.who.int/impact/FinalBrochureWHA2008a.pdf>.
[28] Preemption in Medical Device Litigation: What has changed since *Riegal?*, 28, July 2012, <http://www.semmes.com/publications_archive/litigation/pdf/preemption-in-medical-device-litigation.pdf>.
[29] The Medical Device Industry in the United States,

<http://selectusa.commerce.gov/industry-snapshots/medical-device-industry-united-states>.
[30] Top 40 Medical Device Companies, <http://www.mddionline.com/article/top-40-medical-device-companies>.
[31] Hacking Medical Devices for Fun and Insulin: Breaking the Human SCADA System, 2011, <https://media.blackhat.com/bh-us-11/Radcliffe/BH_US_11_Radcliffe_Hacking_Medical_Devices_WP.pdf>.
[32] Health Care Equipment Could be Vulnerable to Hackers, 22, October 2014, <http://www.thefiscaltimes.com/2014/10/22/Health-Care-Equipment-Could-Be-Vulnerable-Hackers>.
[33] State Economic Impact of the Medical Technology Industry, Lewin Group, 7, June 2010, <http://www.lewin.com/~/media/lewin/site_sections/publications/stateeconomicimpactofthemedicaltechnologyindustry61510.pdf>
[34] Stengel v Medtronic, 13, January 2013, <http://cdn.ca9.uscourts.gov/datastore/opinions/2013/01/10/10-17755.pdf>
[35] Medical Devices: Parallel Claims Against Device Manufacturers post-Riegel?, Beth Rose, 8, August 2014, <http://www.natlawreview.com/article/medical-devices-parallel-claims-against-device-manufacturers-post-riegel>