

Securing Legacy Mobile Medical Devices

Vahab Pournaghshband, Majid Sarrafzadeh, and Peter Reiher

Computer Science Department
University of California, Los Angeles
{vahab,majid,reiher}@cs.ucla.edu

Abstract

Millions of people use mobile medical devices—more every day. But our understanding of device security and privacy for such devices is incomplete. Man-in-the-middle attacks can be performed on typical Bluetooth-enabled mobile medical devices, compromising the privacy and safety of patients. In response, we developed the *Personal Security Device*, a portable device to improve security for mobile medical systems. This device requires no changes to either the medical device or its monitoring software, and offers protection for millions of existing devices. We evaluate our defense mechanism to show that it adds insignificant overhead and analyze its robustness against various attacks.

Key words: medical device security, man-in-the-middle attack

1 Introduction

Studies show that by 2015, over 500 million people will be using mobile health applications [1]. There were approximately 245,000 insulin pump users in 2005, and the market for insulin pumps is expected to grow at a rate of 9% from 2009 to 2016 [3]. Hanna et al. reports that in the U.S. alone there are 25 million people with wireless implantable medical devices (IMD), and about 300,000 of these IMDs are implanted every year [10].

In 2003, the U.S. Food and Drug Administration (FDA) approved a Bluetooth-enabled medical device for the first time [14]. Since then, dozens of such devices have been introduced to the U.S. market for uses ranging from life-sustaining to life-supporting.

While the need for secure mobile medical systems is widely recognized [9, 4, 11, 13], many manufacturers have not addressed the security risks of such devices, and thus have provided little security for either the devices themselves, or for the data they create and transmit.

Communications security is one critical aspect of protecting these devices. Mobile medical devices typically communicate to an intermediate computer that forwards its signals to a healthcare facility. Since such devices are typically used

¹ This work is supported by NSF grant CNS-1116371.

with little or no configuration by a user or healthcare provider, there is ample opportunity for attackers to mislead the device into communicating with a hacker’s machine instead of its intended intermediary. The communication between the device and its intermediary (real or malicious) typically is wireless, making it more susceptible to eavesdropping and injections.

The consequences of attacks can be extreme, potentially allowing attackers to cause the devices to operate in a life-threatening manner. As an example, consider a heart rate monitor carried by a patient that communicates via Bluetooth to the patient’s home computer, which in turn, forwards heart rate data to the patient’s doctor in real time. If an attacker can alter the data to fake a heart attack, the doctor may institute unnecessary emergency measures. Even worse, if the attacker conceals the actual signs of an impending heart attack, the doctor will be unaware of the need for immediate action.

In this paper we demonstrate a successful man-in-the-middle (MITM) attack and its consequences on a commercially deployed pulse oximeter system; we then propose a defense approach against this and similar attacks. After discussing potential security and privacy failures that can result from an MITM attack, we demonstrate such an attack, showing that the device discloses sensitive information unencrypted. This attack shows that these Bluetooth-enabled mobile medical devices can be made to communicate with an unauthenticated intermediary. This attack can be performed by an unauthorized party equipped with a Bluetooth-enabled laptop.

Our study examines the Nonin Onyx II 9550 fingertip pulse oximeter, a typical Bluetooth mobile medical device introduced to the U.S. market in 2008. It measures pulse rate and blood oxygen saturation levels continuously or on demand, and communicates with an access point (AP) to pass this data at a range of several meters. With only the user’s manual and some publicly available information, we were able to launch a successful MITM attack. Although our experiment used the pulse oximeter, the attacks presented can be performed on other devices, such as the A&D Medical UA-767PBT blood pressure monitor, with little modification.

Our approach to reducing this risk does not require rebuilding or altering legacy devices. We propose a personal area network security device designed to interoperate with mobile medical devices. This security device recognizes the security properties and risks associated with a particular patient’s existing devices, and takes measures to lower those risks. Our defense solution works with existing devices and requires no modification to either the device or the monitoring software installed on the AP; it also offloads security from the medical device, reserving the medical device’s resources for only medical functions. Our proposed defense mechanism is designed for generality and wide applicability for this class of medical devices.

The organization of this paper is as follows: Section 2 presents related work, followed by an overview of Bluetooth-enabled medical devices in Section 3. Our attack assumptions and threat model, active MITM attacks, defending against

these attacks, and evaluation are in Sections 4, 5, 6, and 7 respectively. Future work is presented in Section 8, and Section 9 concludes this paper.

2 Related Work

The security of mobile medical devices is a generally recognized problem that has received special attention in recent years [9, 4, 11, 13]. As a result, there has been some work on demonstrating attacks against various mobile medical devices [8, 12]. As complementary research, some work focused on implementing or recommending defensive approaches against these kinds of attacks [7, 18, 15, 8, 12]. Among the proposed defense approaches, IMDGuard [18], Amulet [17], and Shield [7] are three defense mechanisms against attacks on mobile medical devices that require a special-purpose third-party device to facilitate security. Also, Denning et al. [6] propose a class of devices called communication cloakers that would share secret keys with an IMD and act as a third-party mediator in the IMD's communications with external programmers.

IMDGuard proposes changes in the design of future IMDs for a more secure system and does not work with legacy devices. Also, Amulet, by definition, does not work with existing devices since it requires changes to the existing mHealth system including the medical device. For example, it requires the medical sensor to verify that it is indeed the right Amulet before connecting. Shield, however, is the only solution that is designed to work with existing and even already implanted IMDs by requiring no changes to the device. Shield protects an IMD by jamming its IMD messages, preventing others from decoding them, while the authorized intermediary is able to decode them. It also jams unauthorized commands to protect the patient. However, the idea behind Shield may not be applicable to many mobile medical devices that operate on widely used radio technologies such as Bluetooth or 802.11; this is due to both the nature of the radio technology and the potential legal issues of jamming their signals.

3 Bluetooth-Enabled Mobile Medical Devices Overview

In our work emphasis is placed on a class of Bluetooth-enabled mobile medical devices that communicate with an AP. An AP is a cellphone, a home PC or a hospital monitoring system. This is a broad class of medical devices with common characteristics. In a common Bluetooth authentication mechanism, a predefined PIN is required for pairing the two parties. These devices are usually configurable by the AP. Depending on the device, the AP can set a wide range of parameters on the device, from changing date and time, frequency rate and data format, to setting specific therapy management. Moreover, these devices may store the patient's identity information. Normally, some proprietary software is installed on the AP that organizes and visualizes the data, and may report it to the patient and doctor for therapy management.

4 Attack Assumptions and Threat Model

4.1 Attack Assumptions

We first assume that the PIN used in standard Bluetooth pairing is known to the attacker. This is not an ambitious assumption since it is known that the PIN can be deduced by carefully observing the Bluetooth pairing process [16]. However, alternatively for some devices, there are even easier ways to figure out the PIN than exploiting the existing vulnerability in the Bluetooth pairing process. For instance, the pulse oximeter’s static PIN is included in its advertised service name, hence, making it publicly available (“Nonin_Inc.XXXXXXX” where the X’s indicate its six-digit PIN). In some other devices, such as the the blood pressure monitor, a common default PIN is used for all shipped units and is available in a publicly disclosed specification. Secondly, we assume that the attacker is in the proper range to launch the attack (up to 10 m in this case), meaning that it can communicate with both the AP and the device via Bluetooth. Finally, we assume that the attacker knows the type and model of the device the patient is using.

4.2 Threat Model

In this section we enumerate the possible attacks on mobile medical devices that can be leveraged from MITM attacks.

1. Confidentiality: An MITM eavesdropper listens to the communication between the device and the AP. An attacker can also retrieve private identity information by sending bogus requests to the device on behalf of the AP.

2. Integrity: MITM attackers can modify data packets sent by the device to the AP, thus misleading the AP with false data. They can also perform replay attacks and generate fake data or commands.

3. Availability: Attackers can interrupt the communication by simply refusing to pass the data through. More cleverly, the attacker can send unauthorized configuration commands to the device to either keep the device in a state of elevated energy consumption (e.g., by setting it to a higher data transmission rate) or disrupt the connection establishment process (e.g., by changing the PIN).

5 Active Man-in-the-Middle Attacks on Wireless Links

In this section we discuss the technical details of how to perform MITM attacks on Bluetooth-enabled mobile medical devices. We first need to successfully position the man-in-the-middle, and here, we discuss the steps to do it: (1) Jamming Bluetooth: the first step is to force the existing connection to be dropped, making both the device and the AP discoverable and available to pair up. (2) Pairing with the device: the attacker’s machine pairs itself with the device, providing the correct PIN. (3) Pairing with the AP: to deceive the AP into pairing with

the attacker’s machine, that machine needs to masquerade exactly as the device. Hence, the attacker’s machine should advertise the identical service name and available services, as well as using the same PIN as the one used in the device and spoofing the device’s MAC address.

Once MITM is in place, then the attacker can perform the attacks described listed in Section 4.2, given that the protocol used in the device is understood. Since the entire protocol for the pulse oximeter is not publicly available, before performing the attacks we had to deduce the necessary information about it by using reverse-engineering, as described below.

1. Transmission from the Device: After capturing and inspecting the Bluetooth transmissions from the mobile medical device, we discovered the key aspects of the device’s protocol and the data that it sends to the AP. Our analysis of the captured traffic revealed some useful information. We observed that the data is sent from the device as 4-byte long packets. The second and third bytes are pulse rate and oxygen level, respectively. Additionally, we conjectured that the first and last byte are indicating some sort of status. This is because their values rarely change, and when they change, they seem to be independent of data or its fluctuation (Fig. 1). Similar analysis allowed us to reverse-engineer data coming from the blood pressure monitor.




Fig. 1. Proposed format of the communication packet in the pulse oximeter

2. Transmission from the AP: For reverse-engineering packets coming from the AP, we issued commands from the device with different settings and looked at the packet generated from the software. In doing so we observed that the structure of the packet and the contents would only change with packets containing variables (e.g., setting date and time), allowing us to replay previously sent commands. We also learned that there were no packet-specific fields in the packet, such as checksum, packet length or timestamps.

6 Defending Against MITM Attacks

We begin this section with enumerating the characteristics of a desirable defense solution. We then present our proposed solution for preventing such attacks, along with assumptions underlying our solution.

6.1 Desirable Characteristics of a Defense Mechanism

1. Security vs. Responsiveness: An effective mechanism for security should not introduce a significant increase (in medical terms) in the transmission time.

2. Security vs. Availability: A robust defense mechanism should not decrease the functionality of the system. Also, it should not provide new avenues for an unauthorized person to drain a device’s battery. Furthermore, the mechanism itself should not introduce significant power or memory requirements that threaten the availability of the device itself.

3. No Changes to the Medical Device: To secure existing medical devices, the defense mechanism should not require any changes to the device.

4. No Changes to the Monitoring Software: The defense mechanism should not require changes to the implementation of the proprietary monitoring software running on the AP. This, coupled with the previous requirement for “no changes to the medical device,” would improve the security of the existing systems. Note that minor changes to the operating system running on the AP are still acceptable.

6.2 Personal Security Device

In our solution, we propose that a separate wireless mobile device augments the security of mobile medical device systems. As envisioned, this Personal Security Device (PSD) will be small, portable, inexpensive, and easy to use. It can be small enough to clip on a belt or fit in a pocket. The PSD would work with other wireless medical devices to enhance their security and monitor their environment.

The PSD is aware of the suite of wireless mobile medical devices used by the owner, and it has a built-in knowledge of their security properties and vulnerabilities. The PSD takes steps to augment the security of the owner’s devices, such as adding authentication and encryption to data streams.

The PSD can be used as an overlay, changing the transmission path from device→AP to device→PSD→AP (Fig. 2). In this case, even though the PSD could secure the link to the AP, the link from the device to the PSD remains unsecured. This is because we are constrained by not changing the device. In Section 7.2 we discuss possible improvements on the security of this link.

Even if the PSD cannot ask the device to transmit its data stream through the PSD, it may be able to improve the security of the system. For example, if the medical device always pairs with any device knowing its protocol, service name, and PIN, the PSD cannot prevent this device from pairing with others, but it can listen to the signals sent by the device. The PSD can create a parallel data stream containing authentication information (signed by a secret key known only to the PSD) to vouch for the data stream to other machines further along in the data flow of the overall system. This way, integrity attacks leveraged from MITM attacks are substantially harder for attackers to achieve without being detected.

For our PSD system to work, some assumptions must hold. We assume that the medical device pairs with only one AP at a time. Note that in the absence of this assumption, and since we cannot make any changes to the device, we cannot prevent the device from connecting to the attacker’s machine. Moreover, we assume that the device and the AP are not compromised, meaning that the attacker has not been able to alter the hardware or software of the device or

the AP, so they behave exactly as specified. We further assume that the PSD is located close to the device. And finally, we assume that the adversary does not physically try to remove the PSD, or damage it, or remotely hack it.




Fig. 2. Illustration of the mobile medical system when using the PSD

6.3 Required Changes at the AP

The AP needs to understand the new authentication and augmented encryption to be able to communicate securely with the PSD. On the other hand, as discussed in Section 6.1, we want to be able to use the monitoring software as it is, without imposing any changes. Therefore, the AP needs to ensure that the monitoring software still sends and receives the data unencrypted. To accomplish this we installed a virtual machine on the AP and installed the monitoring software on the guest operating system. The host OS, on the other hand, is basically used as a gateway for all traffic coming from or going to the monitoring software, leaving all authentication and encryption to be done on the host OS. For the monitoring software to communicate transparently with the device, we manually created two serial COM ports on the guest OS to emulate both incoming and outgoing Bluetooth communication. Accordingly, we configured two pipes on the virtual machine to redirect unencrypted traffic to and from the serial COM ports.

Our proposed changes at the AP rely on the fact that most modern operating systems either have native facilities for supporting virtual machines of the style we used, or can be easily augmented with other software to do so.

6.4 Discussion

1. In our proposed defense mechanism, we left major security components out of the medical device by introducing the PSD. While this decision was made to secure existing devices by requiring no change to their current design, there are also benefits to this design choice that make it attractive, even for designing other devices in the future:

- a) Mobile medical devices have the unique property that they must fail-open when unbounded access is needed in emergency situations. Simply put, security in life-critical medical devices should never come ahead of accessibility. For instance, if a patient with an implantable defibrillator collapses, the treating doctor would need to be able to communicate with the device to retrieve the patient's information and history and issue necessary commands for treatment. Denying access to the doctor in such a situation is unacceptable. For life-critical medical devices, our defense mechanism approach

complies with the fail-open property since unbounded access to the device can be always granted by simply turning off the PSD.

- b) Resource constraints such as limited memory and battery pose a challenge for security implementation in mobile medical devices. Leaving expensive cryptographic computations to another device would make the device resources more available to life-sustaining functionalities.
2. Our approach requires a less detailed understanding of the device and its protocols, limited to certain security issues, than required by others [17].
 3. Our defense solution makes it possible for the PSD-AP link to use any radio technology other than Bluetooth if desirable. With this, the PSD and the AP could agree on using a different radio technology that is more suitable for that particular environment.
 4. Although we presented this defense solution only on Bluetooth, the idea can be extended to medical devices of other radio technologies. In order to implement this for another radio technology, on one end the PSD needs to be equipped with that particular radio technology capability to connect to the device. On the other end, at the AP, the host OS needs to virtualize the radio interface so that it communicates with the monitoring software running on the guest OS via a pipe. In Bluetooth, as described earlier, we have accomplished this by creating serial COM ports emulating a Bluetooth connection. For 802.11 for instance, one can perhaps modify the implementation of wireless card virtualization so that it communicates via a pipe created in the host OS rather than the actual wireless card on the AP. Chandra et al. [5] provides an implementation of virtualizing a single wireless card.

7 Evaluation

7.1 Performance

Our implementation introduced 783 ± 136 ms delay for every data packet sent by the device and received by the monitoring software when we used a Python implementation [2] of the 128-bit key AES encryption algorithm. This delay is insignificant, even in life-critical medical systems.

7.2 Robustness Analysis

1. Security of PSD-AP link: Since we have complete control over this link, given the availability of resources, we can make it arbitrarily secure by using a strong authentication and encryption, as well as an entirely different radio technology other than Bluetooth, if it is more suitable for our environment.

Bluetooth jamming is one source of denial of service attacks on this link. Even though the PSD does not protect against this attack, this attack is easily detected. Alternatively, the PSD and the AP can switch to using a radio technology that remains unaffected by Bluetooth jamming.

2. Security of Device-PSD link: This link is arguably vulnerable to MITM attacks. This is because an attacker can potentially perform MITM attacks on the link between the PSD and the device, making the entire security system introduced in this paper ineffective. The fundamental challenge to securing this link is that we cannot implement an additional security that requires altering or rebuilding the device. Here we present some recommendations that, while not completely eliminating the possibility of an attack, represent a substantial improvement in minimizing the risk.

- a) *Configuring the device to low power transmission:* If it is feasible for the transmit power of the medical device to be set very low, then it could only communicate with devices that are very close to the medical device, perhaps only to those worn by the patient—namely the PSD.
- b) *Designing an alert PSD:* Unlike the device and the monitoring software, the PSD is designed for security. Hence, an alert PSD would watch for signs of attacks such as MITM and other suspicious events, and would raise an alarm accordingly.

8 Future Work

The PSD idea could be developed into either a self-contained, specialized device, or into a smart phone as an application. Having it as a special-purpose piece of hardware theoretically has some advantages over the smart phone application idea: it makes the system harder to hack, it is less battery-consumptive, and there is no resource contention with other applications. On the other hand, having it implemented as a smart phone application makes it more convenient and readily available to the user. Future work could involve careful investigation of advantages and disadvantages of either option.

We designed the PSD so that it prevents MITM attacks. The other element of the PSD's behavior could be to observe the local environment for signs of attacks on its devices. For example, if the personal security device observes improper attempts by unauthorized devices to pair with the medical device over Bluetooth, it can raise an alert that there is a heightened risk of man-in-the-middle or data stream alteration attacks. A further study could look into feasibility of extending the current PSD to employ more sophisticated security tools, such as medical telemetry anomaly detection and detecting, not just preventing, man-in-the-middle attacks.

Work is needed to find ways to make the use of the PSD acceptable to the class of users it would most benefit.

9 Conclusion

This paper addresses the problem of communication security and privacy for the class of mobile medical devices that communicate via Bluetooth. We presented

the steps we used to launch an MITM attack on such devices. Then, we introduced our Personal Security Device, a separate wireless device that augments security to the existing mobile medical devices to defend against MITM attacks.

References

1. Mobile Health Apps: What Do You Use? <http://www.cbc.ca/news/pointofview/2010/12/mobile-health-apps-what-do-you-use.html>.
2. PyCrypto: The Python Cryptography Toolkit. <http://www.pycrypto.org>.
3. Insulin pumps - global pipeline analysis, opportunity assessment and market forecasts to 2016, globaldata. *Global Data* (2010).
4. AVANCHA, S., BAXI, A., AND KOTZ, D. Privacy in mobile technology for personal healthcare. *ACM Computing Surveys* (2012).
5. CHANDRA, R., AND BAHL, P. Multinet: Connecting to multiple ieee 802.11 networks using a single wireless card. In *INFOCOM. 23rd Annual Joint Conference of the IEEE Computer and Communications Societies* (2004), vol. 2, pp. 882–893.
6. DENNING, T., FU, K., AND KOHNO, T. Absence makes the heart grow fonder: New directions for implantable medical device security. In *Proceedings of the 3rd Conference on Hot topics in security* (2008), USENIX Association, p. 5.
7. GOLLAKOTA, S., HASSANIEH, H., RANSFORD, B., KATABI, D., AND FU, K. They can hear your heartbeats: non-invasive security for implantable medical devices. In *Proc. of the ACM SIGCOMM Conference* (New York, NY, USA, 2011), pp. 2–13.
8. HALPERIN, D., HEYDT-BENJAMIN, T., RANSFORD, B., CLARK, S., DEFEND, B., MORGAN, W., FU, K., KOHNO, T., AND MAISEL, W. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *Security and Privacy. IEEE Symposium on* (2008), IEEE, pp. 129–142.
9. HALPERIN, D., KOHNO, T., HEYDT-BENJAMIN, T., FU, K., AND MAISEL, W. Security and privacy for implantable medical devices. *Pervasive Computing* (2008).
10. HANNA, K. *Innovation and invention in medical devices: workshop summary*. National Academies Press, 2001.
11. KOTZ, D. A threat taxonomy for mHealth privacy. In *Communication Systems and Networks (COMSNETS), 3rd International Conference on* (2011), IEEE, pp. 1–6.
12. LI, C., RAGHUNATHAN, A., AND JHA, N. Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system. In *e-Health Networking Applications and Services, 13th IEEE International Conference on* (2011), pp. 150–156.
13. MAISEL, W. Safety issues involving medical devices. *JAMA: the journal of the American Medical Association* 294, 8 (2005), 955–958.
14. OTT, L. The evolution of Bluetooth in wireless medical devices. *Socket Mobile, Inc. White Papers* (2010).
15. RASMUSSEN, K., CASTELLUCCIA, C., HEYDT-BENJAMIN, T., AND CAPKUN, S. Proximity-based access control for implantable medical devices. In *Proc. of 16th ACM Conference on Computer and Communications security* (2009).
16. SHAKED, Y., AND WOOL, A. Cracking the Bluetooth PIN. In *Proc. of 3rd International Conference on Mobile systems, applications and services* (2005).
17. SORBER, J., SHIN, M., PETERSON, R., CORNELIUS, C., MARE, S., PRASAD, A., MAROIS, Z., SMITHAYER, E., AND KOTZ, D. An amulet for trustworthy wearable mHealth. In *HotMobile* (New York, NY, USA, 2012), ACM, pp. 7:1–7:6.
18. XU, F., QIN, Z., TAN, C., WANG, B., AND LI, Q. IMDguard: Securing implantable medical devices with the external wearable guardian. In *INFOCOM* (2011).