Proximity-based Access Control for Implantable Medical Devices

Kasper B. Rasmussen
Dept. of Comp. Science
ETH Zurich
8092 Zurich, Switzerland
kasperr@inf.ethz.ch

Claude Castelluccia
INRIA
655, avenue de l'Europe
38334 Cedex, France
ccastel@inrialpes.fr

Thomas Heydt-Benjamin
Dept. of Comp. Science
ETH Zurich
8092 Zurich, Switzerland
tshb@cryptocracy.net

Srdjan Capkun
Dept. of Comp. Science
ETH Zurich
8092 Zurich, Switzerland
capkuns@inf.ethz.ch

ABSTRACT

We propose a proximity-based access control scheme for implantable medical devices (IMDs). Our scheme is based on ultrasonic distance-bounding and enables an implanted medical device to grant access to its resources only to those devices that are in its close proximity. We demonstrate the feasibility of our approach through tests in an emulated patient environment. We show that, although implanted, IMDs can successfully verify the proximity of other devices with high accuracy. We propose a set of protocols that support our scheme, analyze their security in detail and discuss possible extensions. We make new observations about the security of implementations of ultrasonic distance-bounding protocols. Finally, we discuss the integration of our scheme with existing IMD devices and with their existing security measures.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design —*Wireless communication*; C.3 [**Computer Systems Organization**]: Special-Purpose And Application-Based Systems—*Real-time and embedded systems.*; J.3 [**Computer Applications**]: Life And Medical Sciences—*Medical information systems*

General Terms

Security, Measurement, Human Factors

Keywords

Distance Bounding, Secure Pairing, Ultrasonic Communication, Access Control, Medical Devices

1. INTRODUCTION

In order to facilitate communication and data readout, new generations of Implantable Medical Devices (IMDs), such as pacemakers, are equipped with radio transceivers. Such interfaces makes it convenient for medical professionals to get access to the data they need but they also introduce some unique security and privacy challenges, access to personal data and the unauthorized modification of IMD parameters being the most prominent [11, 17, 9, 12].

CCS'09, November 9–13, 2009, Chicago, Illinois, USA.

In order to prevent unauthorized access to IMDs, conventional solutions, based on public-key cryptography or on preloaded secret keys cannot be directly applied since they typically also prevent access to IMDs in emergency scenarios where the IMD needs to be accessible to emergency ambulance staff [10]. IMDs therefore create a tension between access control, i.e., patient privacy, and patient safety. Several solutions were proposed to address this problem. Some solutions are based on external devices such as access tokens and communication cloakers [22, 6], whereas others rely on close-range communication channels (e.g., RFID) [15]. In addition to possibly being stolen, lost, or simply forgotten by the patient, external devices also serve as a constant reminder to the patient of her/his condition. Access control solutions based on close-range communication have the advantage of being simple and intuitive, but do not provide any firm guarantees about the range of communication. Namely, an attacker with a strong enough transmitter and a high-gain antenna will still be able to communicate with the IMD even from far outside the intended range (for RFID-based solutions from up to ten meters away [8]). Currently deployed solutions based on magnetic switches are equally based on close-range communication; in addition they do not require any form of authentication to unlock access to the device and are thus inherently insecure (incidents were reported when pacemakers were unlocked by a magnetic field from a patient's headphones [7]).

In this work, we propose a new access control mechanism for implantable medical devices. This mechanism is based on ultrasonic distance-bounding and enables an implanted medical device to grant access to its resources only to those devices that are in its close proximity. Our solution resembles close-range communication solutions proposed in prior work in that it requires a device to be close to the IMD to get access, but differs in that it prevents the attacker from accessing the IMD from further away, regardless of the type of transceiver or antenna he has. Its security relies on the speed of the sound which can not be altered. Moreover, unlike prior proposals, our solution enables IMDs to predefine an exact range from which they can be accessed (with a high degree of accuracy). We achieve this with a new proximity-based device pairing protocol based on ultrasonic distance bounding. In this protocol, messages are cryptographically tied to the distance bounds measured by the IMD, to the device that requests access. We analyze the security of our protocol in detail and discuss possible extensions related to efficiency and DoS protection.

We demonstrate the feasibility of our approach through experiments in an emulated patient environment. We show that, although implanted, IMDs can successfully verify the proximity of other devices. We further make new observations about the security of implementations of ultrasonic distance-bounding protocols. We show that without appropriate shielding of their ultrasonic

circuitry, implementations of these protocols are vulnerable to attacks resulting from induced current in the audio receiver circuitry. We further show that given that our solution relies on ultrasonic distance-bounding, it can be implemented at low cost on existing IMD platforms (note that some pacemakers already come equipped with speakers).

Finally, we discuss the integration of our scheme with other solutions proposed for access to IMDs. We show that our solution can be easily combined with solutions based on security credentials or tokens as well as with solutions that aim to prevent battery draining Denial-of-Service attacks on IMDs (e.g., [21, 9]). It also naturally integrates with solutions based on sonic user alerts [11].

We summarize our contributions in the following points:

1. We propose a new access control mechanism for implantable medical devices that enables IMDs to predefine an exact range from which they can be accessed.

2. We demonstrate the feasibility of ultrasonic distance bounding between implanted medical devices and external readers using our prototype implementation.

3. We show that without appropriate shielding all devices using ultrasonic distance bounding are vulnerable to compromise by inducing signals in the ultrasonic circuitry.

4. We show that our solution can be easily combined with existing solutions and implemented on existing platforms.

The rest of the paper is organized as follows. In Section 2 we give a detailed description of the problem, in Section 3 we present our access control scheme based on a proximity aware device pairing protocol. In Section 5 we provide the details of the implementation of our prototype, in Section 4 we describe extensions to the main protocol and in Section 6 we cover related work. We conclude the paper in Section 7.

2. SYSTEM AND ATTACKER MODEL

In this section, we present our system and attacker models.

2.1 System Model

In this paper we focus on access control, specifically in the context of implantable medical devices. Access control in this context means that a reader (potentially malicious) will try to gain access to an implantable medical device in order to readout data or send commands. The reader can be either a handheld unit or part of a bigger system but the assumption is that it is not subject to tight power and/or computational constraints. The medical device can be any device implanted into the human body—including pacemakers, implantable cardiac defibrillators (ICDs), drug delivery systems, and neurostimulators.

Implantable medical devices are used to help manage a broad range of ailments such as cardiac arrhythmia, diabetes and Parkinson's disease. They are implanted 2-3cm below the skin and electrodes are then connected to whatever organ that needs monitoring, e.g., the heart in the case of a pacemaker.

These devices do not have wired interfaces (e.g., a USB interface), screens, keyboards or other peripherals that can be used to enable access control; instead, these devices only rely on wireless interfaces. The reason for having a radio interface in an IMD is that a doctor or medical professional can interact with the device quickly and easily. This is not only useful during normal consultations with physicians but is also relevant in emergency scenarios where emergency personnel (possibly in a different country) need access to the patient's IMD. That means that a device must be accessible in the noisy and dynamic environment of a moving ambulance and at the same time prevent unauthorized access to potentially sensitive medical data.

We consider an IMD that can operate in two different modes. In *normal mode* a reader needs to be in possession of a shared key in order to talk to the IMD and in *emergency mode* a reader just needs to be within a certain security range. In other words the emergency mode relies on proximity alone to authorize a reader.

2.2 Attacker Model

We consider two different attack scenarios. In the first scenario, the attacker wants to get access to medical data stored in the implantable device or change device settings [12]. The motivations for this kind of attack can be anything from identity theft or blackmail to simple curiosity or targeted advertising.

In the second attack scenario we consider an attacker that does not care about establishing a connection with the implantable medical device but instead wants to impersonate a device and make a reader talk to him. This attack might be executed by someone who wants to prevent care in an emergency situation or it could be performed by the patient himself for the purpose of insurance fraud.

We place almost no restrictions on the attacker in terms of communication, i.e., he can send and receive arbitrary radio and audio signals. We do assume that the attacker is subject to common computational bounds, i.e., he is not able to reverse one-way functions or solve the discrete logarithm problem. We also assume that the attacker is outside the security range defined in the IMD (typically <10cm). If the malicious reader is inside the security range and the IMD is in emergency mode, the reader has free access by design. The size of the security range is discussed in Section 4.

Because the implantable medical devices run on batteries they are naturally energy constrained. That makes energy draining and DoS attacks a danger to IMDs. Our protocols are designed with energy conservation in mind, however, such attacks are not specifically addressed in this paper. We focus on attacks on the key agreement and proximity features of our protocol but it is worth noting that our scheme nicely integrates with existing solutions to protect against DoS/Energy draining attacks [21]. Some of these solutions are discussed in related work in Section 6. We also do not specifically address attacks on patient privacy in which the attacker tries to check if the patient is wearing a pacemaker [9].

3. PROXIMITY-BASED ACCESS CONTROL FOR IMPLANTABLE MEDICAL DEVICES

In this section we will describe our proximity-based access control scheme for implantable medical devices. In our scheme the access control is based on device pairing. In order for a reader to talk to an IMD it must first run a device pairing protocol and generate a shared key. This shared key is then used to gain access to the device, either to send it commands or to readout medical data.

The core of the scheme, namely the proximity aware device pairing protocol between a hand held reader and an implanted medical device will be presented in this section. Extensions to the protocol are presented in Section 4. The protocol uses ultrasonic distance bounding to determine the distance between the reader and the device. As is common practice, we will use the terminology *prover* and *verifier* to denote the two parties throughout the rest of the paper. The prover is the reader that must prove its proximity in order for data transfer to commence. The verifier is the implanted medical device that must verify the distance to the prover before accepting the connection.

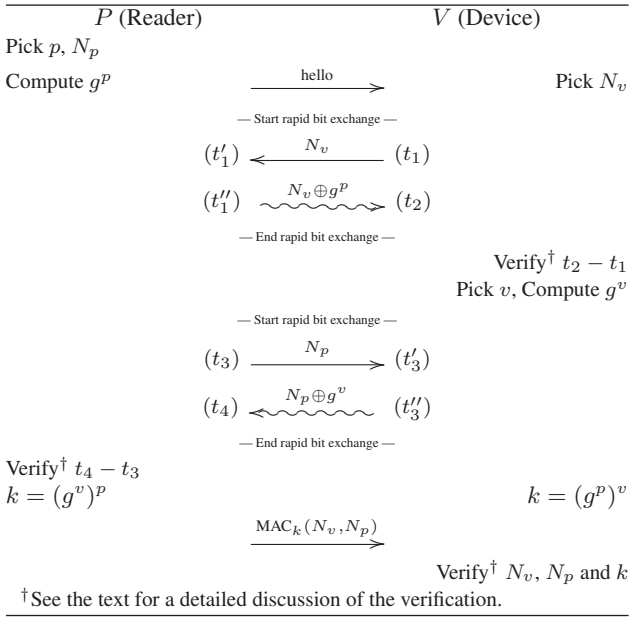

Figure 1: Device pairing protocol. If this protocol is successfully executed both parties know that the key contribution from the other party came from within a distance determined by $t_2 - t_1$ ($t_4 - t_3$ respectively). The prover and verifier also shares a secret key k.

3.1 Protocol Description

The device pairing protocol is shown in Figure 1. The prover will first pick a secret exponent p and a nonce N_p and then compute the public DH contribution g^p. These computations are done in advance so they will not interfere with the time-critical distance bounding steps. A 'hello' message is sent by the prover to initiate the protocol. When the verifier receives the 'hello' message it will pick a nonce N_v and begin the rapid bit exchange phase. The verifier will send a single bit of N_v to the prover and record the time of transmission (t_1) so the time-of-flight can later be calculated. The distance bounding phase must be done bit-by-bit to avoid distance shortening attacks [2, 5].

The message containing the first bit of N_v is received by the reader at time t'_1 but given that the reply must be sent via the sound channel and that the speed of sound is relatively slow compared to the propagation speed of the radio message and the delay at the prover, we consider $t_1 = t'_1 = t''_1$. The error resulting from this assumption is negligible as long as the prover replies immediately. This will be described in more detail in Section 3.4.

The prover xor's the single bit message with a single bit of g^p and sends it back as a sound message. The verifier receives the sound message at time t_2. As described above the verifier uses the time difference $t_2 - t_1$ to calculate the (upper bound) distance to the prover. The distance is calculated as $d = v_s(t_2 - t_1)$, where v_s is the speed of sound in meat (approximately 1500m/s). If this distance is less than some predefined value, say, 5cm the protocol continues, otherwise the verifier will terminate the session. After all the bits of N_v and $N_v \oplus g^p$ have been exchanged, and passed the time-verification, the message is accepted and the DH contribution is assumed to originate from a very close reader.

After the prover has verified that the reader is within the required distance, the verifier picks v and computes g^v. A similar distance bounding step, i.e., a rapid bit exchange with radio challenges and response via the sound channel, is then repeated from the verifier to the prover to ensure that the reader is talking to a device in it's proximity. This is needed to prevent a (possibly far away) attacker from impersonating a device.

Finally, in order to let the device know that a key was successfully established, the prover sends a final message to the verifier containing a message authentication code (MAC) of the two nonces N_p and N_v. At this point, the verifier knows that a key has been established and data transfer can continue encrypted.

3.2 Security Analysis

Central to our device pairing protocol is the unforgeable assurance of proximity. That assurance comes from tying the DH key contributions from each party to the distance between them, by transmitting g^v and g^p over the sound channel. In this section we will go through the protocols resilience to attacks from outside the security range. Note that since there is no notion of identity in this protocol the reader is considered authorized if it is within a specified distance, thus the attackers we are looking at here are attacking from further away. We further assume that the attacker cannot send data on the sound channel faster than the speed of sound[1].

One possible attack is for the attacker to guess N_v and then generate the sound messages in advance. If the attacker is able to generate all the sound messages and send them at the appropriate times, the attacker could pretend to be close to the verifier while actually being far away. That means that the nonce N_v must be sufficiently random to make guessing infeasible.

The nonce N_v is sent in the clear since it is the timing of the sound message that proves the proximity of the reader. An attacker who is further away than the allowed distance will receive the nonce at more or less the same time (the propagation time of radio signals is negligible when compared to the speed of sound) but, because he has to wait for N_v before he can create a valid sound message, his sound message will not be able to reach the prover in time, i.e., the prover will be able to measure the distance to the attacker and conclude that he is too far away.

A similar distance bounding step is repeated in the opposite direction. This proves to the reader that the IMD is also within the specified distance, eliminating impersonation attacks. Since the two DH contributions are sent over the sound channel they are directly linked to the distance between the reader and IMD, which also makes the key $k = g^{vp}$ directly linked to the distance as well.

In order to limit the effectiveness of battery draining attacks the IMD only generates its public DH contribution after the distance to the reader has been verified. That way only the initial nonce must be generated at the start of each session.

The final message from the prover to the verifier confirms the key. After executing this protocol the verifier knows that a valid key has been generated with a prover and that this prover is within the allowed distance. At this point the verifier can start transmitting data using the generated key k or send another message to the prover confirming the key.

3.3 Side Channel Attack Protection

One of the most important assumptions in our security analysis is that the attacker cannot send data on the sound channel with a signal that propagates faster than the speed of sound. This is a common assumption and it is made in most, if not all, the sonic and ultrasonic distance bounding protocols in the literature, including [16, 27, 3, 19].

While this assumption sounds perfectly reasonable there are pitfalls that an attacker might utilize. While fine tuning our imple-

[1]In Section 3.3 we discuss some interesting and novel ways a powerful attacker might be able to get around this assumption.

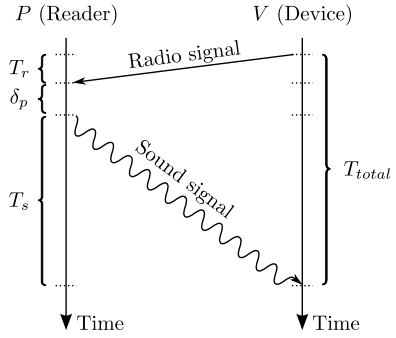

Figure 2: The propagation times and processing delays involved in an ultrasonic distance bounding challenge-response. (Not to scale). T_r is the propagation time of the radio signal, δ_p is the prover's processing delay and T_s is the propagation time of the sound signal.

mentation we found that it is possible to send a radio signal to the IMD that will induce a current in the audio receiver circuit just as if the IMD received a sound signal. This was possible even though the prototype was not designed to receive RF at all, just the fact that there where two small wires going from the reception circuit to the piezo element (working as a microphone) was enough to pick up a radio signal of about the same order of magnitude as the audio transmission.

The only countermeasure to this is effective RF shielding. This is arguably an engineering problem but one that is very easily overlooked. If proper shielding is not in place, a strong attacker can effectively send an 'audio' transmission at the speed of light!

3.4 Propagation Time and Processing Delay

In the description of the protocol we state that the propagation time of the radio signal and the delay at the prover is negligible. More precisely we claim that $t_1 = t_1' = t_1''$. In this section we will show that the error resulting from this assumption makes little to no practical difference.

The potential time-error sources are shown in Figure 2. Our above mentioned assumption corresponds to the following two assumptions $T_r = 0$ and $\delta_p \ll T_s$ and allows the verifier to compute an upper bound on the distance to the prover as $d_{vp} = T_{total} \cdot v_s$, where v_s is the speed of sound in meat. We will look at the two assumptions one at a time.

The true value of T_r is $d_{vp} \cdot c$ where c is the speed of light. For a distance of $30m$ (which is about 100 times the intended operation distance of the protocol) the true value of T_r is $0.1\mu s$ with corresponds to an error in the distance measurement of $34\mu m$ (assuming a speed of sound $340m/s$). That is way below the distance granularity of our measurement setup and we therefore consider it irrelevant.

The delay at the prover δ_p is the time it takes from the prover receives the first bit of the radio signal until the first bit of the ultrasound signal is in the air. The actual delay will vary depending on the speed of the prover's hardware and the details of the implementation. δ_p measurements from our setup are described in Section 5.2; using the average delay we have that $\delta_p = 412ns$. This delay corresponds to an error in the distance measurement of $140\mu m$ (assuming a speed of sound $340m/s$). This error is also way below our distance granularity.

There is one final important point regarding timing. The speed of sound is higher when the sound propagates through the human body than when the sound propagates through air. According to [18] the speed of sound through the human body is approximately 1500m/s which is about three times the speed through air. Because IMDs are implanted inside human tissue we have to assume a speed of sound of $1500m/s$ when defining the maximum distance from which the device can be accessed. Any distance the signal has to travel through air to get to the reader will be counted three times because the signal travels three times slower. That means that any additional distance to an attacker outside the allowed access radius is amplified thus making it even harder to cheat the system.

4. PROTOCOL EXTENSIONS

In this section we will present several extensions to our proximity aware device pairing protocol.

4.1 Combining Proximity and Credential-Based Solutions

As we describe in more detail in Section 6, it is very likely that patients will be provided some form of credential (a smart card, USB stick or password) that shares a secret with the implanted medical device. This credential would be used by the doctor to actually get access the IMD when necessary.

This solution is actually quite attractive since it solves the authorization and authentication issues. In fact, by physically giving the credential to the doctor, the patient is explicitly authorizing the doctor to get access to his IMD. Furthermore since the credential shares a secret with the IMD, it can be used by the reader to get access to the IMD (access control) and bootstrap a key that is used to securely exchange data. A patient might feel safer if the security of his IMD is based on some secret credential in his possession, rather than on a proximity-based solution.

However the credential-based approach has several drawbacks that our scheme can help solve. First, if the credential gets stolen or duplicated, any attacker can get remote access to the IMD. Second, the doctor does not have the insurance that his reader is actually communicating with the patient's IMD. In fact, nothing prevents the patient from borrowing the credential of a friend and have the doctor's reader communicate with the friend's IMD, who is sitting next door. This attack could be, for example, used for Medicare or insurance fraud purposes. Third, if the patient does not carry his credential, no one can access the IMD even in case of emergency. This is clearly not acceptable since this can put the patient's life in danger.

We believe that our scheme can nicely complement the credential-based solutions to solve these three issues. We consider two modes of operation. In the *normal mode of operation*, the patient carries the credential token and provides it to the doctor that needs to access the IMD. In the *emergency mode of operation*, the doctor does not have access to the credential token, either because the patient has lost/forgotten it or the token is out of order.

Normal Mode of Operation

The patient carries an authorization credential token (USB token, smart card, password, etc.) that shares a secret key k_{shared} with the IMD. When a doctor needs to access the IMD, he gets the credential from the patient and provides it to the reader. The same proximity aware device pairing protocol shown in Figure 1 is run between the reader and the IMD except the shared key k_{shared} is included in the MAC in the final message. Once the protocol has been executed, each party has the assurance that the other party is within its security range and has derived a key k that is used to secure their future communication.

By verifying that the IMD is in the proximity of the reader, the doctor has the assurance that his reader is communicating with the patient's IMD.

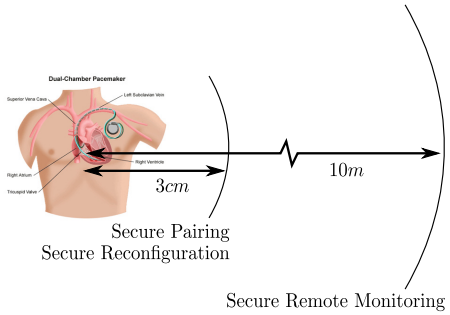

Figure 3: Security Regions. Different types of operations have different security requirements, and therefore security regions. Critical operations, such as IMD reconfiguration, should use a small security region. Monitoring operations, with are not life threatening can use a larger security region, in order to improve usability.

Note that since, in this mode of operation, the IMD and the reader share a secret, the Diffie-Hellman key exchange could easily be avoided if necessary. In fact, both parties could derive a key k from the shared secret k_{shared} and the exchanged nonces. However, the ephemeral Diffie-Hellman key exchange protocol provides forward security, which can be a valuable property.

Emergency Mode of Operation

In this mode of operation, it is assumed that the authorization token is not available. With most existing systems, in this situation, wireless communication is not possible unless the IMD is activated by a magnetic read switch. However, as described in Section 6, these physical backdoors have many drawbacks. We, instead, propose to use the protocol shown in Figure 1 in this mode of operation.

With this solution, both the reader and the IMD verify that they are within each other's security range and generate a temporary secret key. An attacker won't be able to get access to the victim's IMD from a remote location, however, he could potentially establish a key with the IMD if he gets close to the patient, without having to steal his credential. We therefore suggest, that in the Emergency mode of operation, the security range should be much smaller than in the normal mode of operation. We believe that a security range of 2-4 centimetres should be quite appropriate, since this would require the attacker to almost have physical contact with his victim.

Note that inputs from other sensors could be used to reinforce the security of the emergency mode of operation. For example, if the IMD is equipped with an accelerometer, the policy of the IMD might be to verify that the reader is close, as proposed by our scheme, but also that the patient is lying down. This would provide an additional level of protection. Furthermore if the IMD detects an emergency situation (stroke, heart failure, etc.), access control could be deactivated all together. Access control, in this case, is probably not the biggest concern.

4.2 Proximity-Based Commands

Until now, our proximity-based scheme has been proposed to secure the IMD–reader pairing during the normal and emergency modes of operation. However, this approach can be extended to any other aspect of IMD–reader communication.

A doctor might want to access an IMD for several reasons. One reason could be to remotely monitor a patient and retrieve logging/history data. Another reason could be to modify the parameters of the IMD or reconfigure the device.

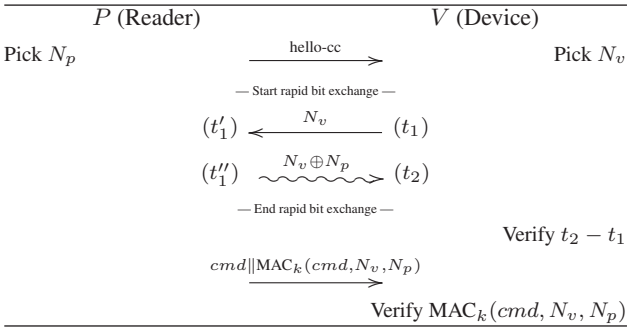

Figure 4: Message proximity verification. With this protocol the verifier V can verify that the command cmd came from the distance defined by $t_2 - t_1$, or closer.

The second type of operation is clearly more critical and requires stronger security, since it can potentially threaten the life of the patient. The first type would only violate privacy if performed by a non-authorized user. It is therefore reasonable to apply different security policies for each of these operations. We propose that as long as the implanted medical device is in the normal mode of operation, critical commands are only processed if issued by a reader that is in its proximity as shown in Figure 3.

This proposal authorizes remote monitoring of an IMD via a secured channel but excludes remote reconfiguration of an IMD. We believe this is a reasonable approach and provides an acceptable security/usability trade-off.

In order to verify the proximity of the reader when it sends a command we propose the command proximity verification protocol, illustrated in Figure 4. It is assumed that the reader and the IMD share a secret key, k, i.e., that both devices have been securely paired already.

When a reader wants to send a critical command to an IMD, it starts by sending a 'hello-cc' to initiate the protocol. The IMD picks a nonce N_v and replies with the first bit of N_v. The IMD also starts a timer so the time-of-flight of the sound message can be measured. The reader responds immediately with a single bit of its own nonce xor'ed with N_v and this continues until there are no more bits in the nonces, or until the IMD aborts the protocol because the estimated distance is outside the security range.

Once the distance bounding phase of the protocol is over the reader sends the command cmd along with a MAC of the command and the nonces. If the IMD is able to verify the MAC it knows that cmd came from within the security distance and will process the command.

Although it is assumed in the proximity-based command protocol that the two devices share a secret key, this protocol could still be useful in scenarios where the only policy for being able to issue command is to be close the device. The modification to the protocol would then be to replace the MAC function with a regular hash function. The security would, of course, be lower but could still be acceptable for some applications.

4.3 Robustness

Because robustness is a key design criterion we propose a method to allow the proximity aware device pairing protocol in Figure 1 to continue, despite transmission errors on the sound channel. This proposal is an optional addition to the protocol and is meant to enable device pairing in extremely loud environments at the cost of some security.

The proposal is, that after the rapid bit exchange phase, the prover (or verifier) sends a radio message containing the exact same data

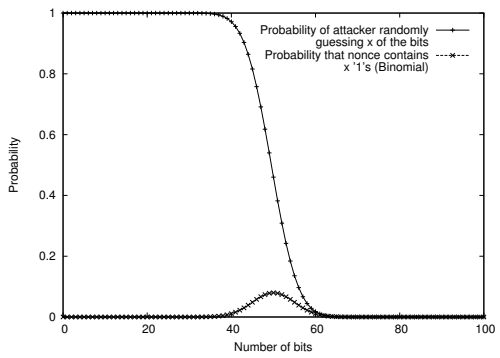

Figure 5: Probability that an attacker can guess x of the bits in a nonce. This assumes that the attacker is guessing each bit randomly with a probability of $1/2$ and applies if the receiver allows bit errors.

$(N_v \oplus N_p)$ as was sent in the sound messages. Doing that will enable the verifier (or prover) to use the arrival time of the sound messages to detect proximity but since the same data was transmitted via the radio channel (which presumably is immune to audio noise) it doesn't matter if part of the audio message is wrong. It should be emphasized that this extra radio message is sent after the distance bounding phase has completed successfully. In order for an attacker to abuse this robustness proposal he must already have cheated the distance bounding phase, i.e., sent all replies at the correct times, otherwise the protocol would have been aborted.

If the verifier (or prover) is willing to accept some transmission errors in the audio messages, it reduces the guessing space for the attacker. However, as long as enough bits are correct, the verifier (or prover) can be fairly certain that the audio messages where not guessed in advance and sent by an attacker. In the following we discuss some guessing strategies that an attacker might use to attack a version of the protocol that allows bit errors on the sound channel.

The optimal guessing strategy depends on how the nonce is generated but if we assume a truly random nonce, most guessing strategies will give the same result, on average. In a guessing strategy where the attacker tries to guess each bit randomly, i.e., '1' with probability $1/2$ and '0' with probability $1/2$, the probability of guessing *exactly* m out of N bits is $P(m) = \binom{N}{m} p^m (1-p)^{N-m}$, assuming each guess is independent. In Figure 5 this binomial distribution is plotted along with the probability of guessing *at least* m out of N bits. This probability is given by the sum of the probability of guessing exactly m bits and the probability of guessing exactly $m+1$ bits, etc.

$$P(\text{at least } m \text{ bits}) = P(m) + P(m+1) + \cdots + P(N)$$
$$= \sum_{i=m}^{N} \binom{N}{i} p^i (1-p)^{N-i}$$

The exact amount of correct bits that the IMD will require will depend on the current mode of operation and the security policy in effect. However, we propose as a reasonable trade-off between security and robustness, to require at least 75% of the bit to be correct. If $N = 100$, that will result in a probability for an attacker to fake his distance of approximately 2^{-22}.

Note that this is an upper bound, since this computation assumes that all 100 sound messages sent by the attacker passed the distance bounding test, i.e., that they were sent at the correct time such that they appear to be from someone in the proximity of the device.

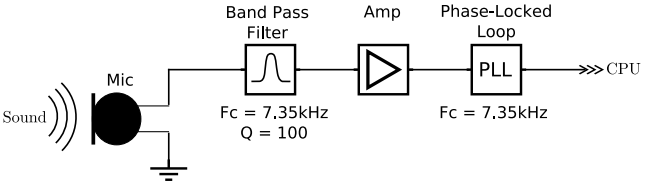

Figure 6: The analog portion of the receiver used in the long distance mode of our proof-of-concept prototype. A highly selective bandpass filter is used to reject environmental noise, then a phase locked loop is used for detection of the communications frequency.

5. PROTOTYPE IMPLEMENTATION AND MEASUREMENTS

In this section we present the implementation details of our prototype and the details of our measurement setup.

5.1 Construction of Proof-Of-Concept Prototypes

In order to test various properties of our system we built proof-of-concept prototypes of both the prover and the verifier, and programmed them with selected portions of the protocols and lower level overhead, sufficient to measure several properties of the system. We describe here the high level technical details of the prototypes.

The prototypes are respectively a prover and a verifier with analog circuitry for RF and sonic communication and ATMega644p microcontrollers running at 20MHz for computation and control. The ATMega644p microcontroller is representative of medium-cost general purpose microcontrollers for embedded systems. The prototypes have two modes of operation: a precision near distance mode, and a long distance mode.

The near distance mode is designed to provide good accuracy in distance measurement, while simultaneously providing a baud rate suitable for fast transaction times. In this mode the analog portion of the receiver consists of a 4 poles VCVS active high pass filter to exclude the majority of environmental noise, and two stages of amplification. We designed the final amplification stage to clip the signal thereby producing a TLL square wave as input to the microcontroller where the frequency detection was implemented. This mode of operation uses a carrier frequency of 13.6kHz and communicates using binary frequency shift keying with a modulation index of 500Hz. We found this setup to be suitable for reliable communication at speeds of up to 1kbaud.

The near distance mode is only intended for the case where the programmer is in very close proximity to the IMD. In order to additionally support the case where it is desirable (perhaps in a lower security setting) to communicate with a programmer at a greater distance, the long distance mode (whose analog receiver is depicted in Figure 6) has a much higher frequency selectivity and performs frequency detection in the analog domain, rather than in the digital domain as in the near distance mode. The analog signal path for the long distance mode consists of a state variable band pass filter with $f_c = 7.35$kHz and a quality factor of $Q = 100$ (very high frequency selectivity). The filter output is amplified and then fed into a phase locked loop timed to detect the f_c of the filter. The phase locked loop outputs TTL level pulses to the microcontroller when the mark frequency is detected, thus supporting the same BFSK communication scheme used by the near distance mode. Long distance mode has less accuracy in distance measurements for reasons discussed below.

We measured the power consumption of the microcontrollers during peak computation and found the receiver to consume $0.15W$ at $5VDC$ and the transmitter to consume $0.17W$ at $5VDC$. We did not attempt to optimize power consumption through turning off unused portions of the controller, or using lower power states during periods of reduced computational demand. The analog portion of the receiver consumes $0.13W$ at $10VDC$. All measurements taken with current sensing laboratory bench power supplies.

Assuming (pessimistically) that authentication takes 1 second, the energy consumed by our prototype would be 0.28J. To place this in context, a defibrillation shock of 10J would be a medium energy therapy for several common conditions, and an Implantable cardioverter defibrillator is expected to be able to provide such a therapy many times. Our prototype can run an authentication protocol around 36 times and consume as much energy as a single such therapy.

5.2 Turn around time

In order for the distance bounding to be sufficiently accurate to provide security it is important that there be only a short lag time between the reception of a bit on the radio channel and the corresponding transmission of the response bit on the sound channel.

We performed measurements by connecting an oscilloscope lead to the microcontroller pins on which demodulated data are received from the radio and sent to the sonic transducer respectively. Using this setup we measured the time required to detect incoming data, store it into memory, calculate the response based on this challenge, and output the response. We found that approximately $412ns$ (mean of 20 trials) are required for these combined operations, with a standard deviation of $45.4ns$. During this amount of time sound travels around one hundredth of a centimeter, which indicates that the digital domain overhead of our protocol has a negligible impact on distance estimation when run on reasonably modern embedded technology.

5.3 Distance Measurement

To test the accuracy of distance estimation we measured the time of flight as the difference between the time one bit of information reaches the output of the transmitter and the time that bit can be recognized by the receiver. This latter time includes the time of flight of the sound, and the delay introduced by the entire analog signal path. The signal thus produced was measured at the point where the input enters the receiver's microcontroller.

In our first series of distance measurement experiments we evaluated the near distance mode's distance measurement accuracy at different distances in air, and through $2cm$ of meat and $1cm$ of air. This latter case represents a likely practical scenario for emergency near distance communication, and therefore requires that the programmer be very close to the medical device. Because of the amount of RF noise present with the transmitter and receiver so close together we performed this series of experiments with the receiver mounted in a Faraday cage. In a real IMD the shielding would be integrated into the housing of the IMD itself. To simulate communications through a patient's abdominal wall we implanted the transmitter in $8kg$ of ground beef with at least $2cm$ of meat surrounding the transmitter on all sides. We implanted the transmitter rather than the receiver in this series of experiments both to test the time of flight calculation available to the programmer device in our protocol, and also because the receiver, being in a Faraday cage, would have been difficult to implant.

Through air we measured every distance between $1cm$ and $100cm$ at $2cm$ intervals with 5 observations at each distance. We found the accuracy to be $\pm 1.78cm$ with a standard deviation of $1.59cm$.

Figure 7: To test time of flight distance calculations we implanted the receiver's microphone inside a portion of the abdominal wall of a cow. Shown here connected to the analog portion of the long distance receiver.

Through $2cm$ of meat and $1cm$ of air we found the near distance mode to be accurate to within $\pm .01cm$ with a standard deviation of $0.5cm$ (mean of 14 trials). This provides suitable accuracy for the IMD's emergency mode of operation (see Section 4).

In our second series of distance measurement experiments, we evaluated the long distance mode's distance measurement accuracy through $3cm$ of meat and additional distances in air as shown in Figure 7. We measured total (air + meat) distances from $10cm$ to $150cm$ at $10cm$ intervals, as well as at $5cm$. We recorded 10 observations at each distance.

The implantation setup for our second series consisted of implantation of the receiver's microphone inside a portion of the abdominal wall of a cow measuring slightly more than $15cm \times 9cm \times 6cm$. The microphone was implanted $3cm$ deep.

Figure 8 shows the results of these measurements. Measured distance includes both air and meat between the sound transmitter and receiver. The estimated distance shown is calculated based on the average speed of sound in air at sea level multiplied by the time of flight adjusted for the constant propagation delay through the filter and PLL frequency detector used by the long distance mode.

These measurements show that the prototype has a precision of $\pm 9cm$ with standard deviation of $3cm$.

We see from these experiments that the long distance mode has a greater range, but poorer accuracy than the short distance mode. The inaccuracy of the long distance mode is largely due to the phase locked loop used for frequency detection. A phase locked loop consists of an oscillator at a multiple of the mark frequency with feedback mechanisms which cause this oscillator to lock on to (synchronize with) an incoming signal. The time required to lock on to a new incoming signal is somewhat non-deterministic as it will depend on many factors including the initial phase difference between input and the PLL's internal oscillator. Fortunately, however, $9cm$ is sufficient accuracy for long distance mode operation in which the reader may operate from a distance of several meters.

6. RELATED WORK

Several approaches have been proposed to solve the IMD security problem described in the previous sections. We briefly review

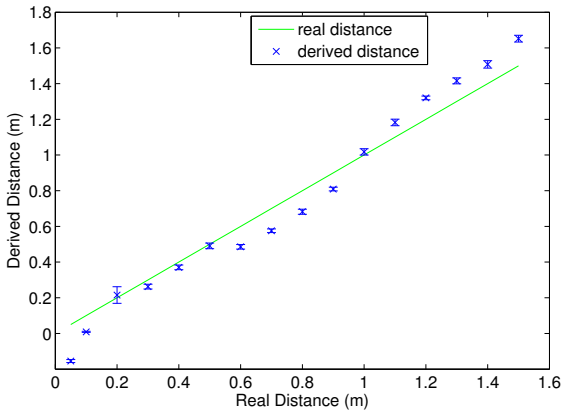

Figure 8: Real distance (constant $3cm$ meat plus variable distance in air) versus distance computed by proof-of-concept prototype in long distance mode based on time of flight. The line $y = x$ is shown for reference. Error bars indicate one standard deviation. The average accuracy is $\pm 9cm$ and average std deviation is $3cm$

them below and we contrast them to our proposed solution where appropriate.

Token-Based Approaches

With token-based approaches [22] the patient gets an access token (e.g., USB stick, smart card) that is configured with a credential (often a secret key) that is shared with the IMD. When a doctor needs to access a patient's IMD, he gets the access token from the patient and provides the credential to the reader. The IMD and the doctor's reader can then establish a secure link that is used to download data or send programming commands. This approach is very practical, however, it has several drawbacks: it does not protect against the loss or theft of the token, it creates a safety problem since the IMD is inaccessible if the user does not carry the token with him and it serves as a reminder to the patient of his health state.

In [6], the authors propose a defensive technique called Communication Cloaker. A Cloaker is an externally worn token, for example a bracelet, that shares a secret key with the IMD. The cloaker broadcasts periodic beacons that are used by the IMD to detect its presence. As long as the cloaker is detected, the IMD will stay mute to any request coming from a reader. If a doctor needs to access the IMD, the patient has to remove his cloaker to allow communication. This approach is interesting and provides an elegant solution to the security/safety tension. However, it suffers from the same drawbacks as the token-based approach described above. Furthermore, it is unclear how this solution copes with jamming attacks that would prevent the cloaker from being heard by the IMD.

Certificate-Based Approaches

With certificate-based approaches [9] the IMD is configured with the public key of a trusted party. When a reader wants to access the IMD, it needs to contact the trusted party and get a valid certificate (credential). This credential is then used to establish a secure link between the IMD and the reader. This approach has several drawbacks. First, it requires the reader to go online and contact the trusted party. This is a very strong requirement since Internet access might not be available everywhere (especially in case of an accident on the road). Second, it requires the deployment of a (global or at least nation-wide) certification authority, which is a costly solution.

User Alerts

In this proposal [11], the IMD emits an alert signal (sound, vibrations, etc.), when it is engaging in wireless communication. This proposal does not control access to the device, but rather alerts the patient about an interaction. This solution might seem simple and attractive but it does not work in noisy environments, it does not (in itself) enable the user to react to the alert and creates new privacy issues (by emitting a sound, the IMD is advertising its presence). As we discussed, this solution can be naturally integrated with our proposal since both proposals rely on the emission of sound and thus require the same hardware.

Proximity Based Access Control Approaches

Some schemes (e.g., [21]) propose to disallow long distance wireless communication with the IMD until a proximity based procedure has completed. For example, in most existing solutions, wireless communication won't be possible unless the IMD is activated by a magnetic switch. A magnetic field is chosen as an input channel since the channel is quite simple, and standardized. Due to the faster drop-off in field strength of a magnet compared to radiant energy the channel implies a certain degree of proximity. In current generation devices, this approach is used, for example, to disable defibrillation shocks from the device while paramedics or ER staff are diagnosing a patient's EKG (which would be interrupted by such a therapy), or to prevent shocks from a device that has run amok. This solution is quite insecure, since the presence of any strong magnetic field will trigger the switch. Incidents were reported where a magnetic field from a patient's headphones tripped the read switch [7]). Since the activation of the read switch usually enables wireless communication, this means that someone with headphones in their pocket might be vulnerable to attacks.

The secure telemetric link solution [1] proposes to use a physical backdoor to verify that the reader is close to the IMD. When the reader wants to access the IMD, it sends an "activation message", over the wireless channel to the IMD that activates the backdoor circuitry. The reader then gets close to the IMD with, for example, a near-field magnetic sensor. If the IMD detects the reader's sensor, it sends the authentication key over the wireless link, using a very low transmission power. The reader then gets the key and can communicate with the IMD via the wireless link. If the IMD does not detect the sensor within a given time frame, it deactivates the backdoor. This solution has all the problems of the magnet-based scheme described previously. In addition, it is not secure against an attacker that uses special equipment (e.g., high-gain antennas) to eavesdrop on the key.

Other schemes [15, 23] also use short-range communication technologies (such as IR, Bluetooth, RFID, etc.) to guarantee proximity. However, all approaches based on short-range communication technologies are vulnerable to attacks since an attacker can easily increase the communication range using powerful and sensitive transceivers and high-gain antennas. These schemes are therefore not secure, since an attacker can eavesdrop and access IMDs from far away. Essentially, with such approaches it is hard to say from which distance access is no longer possible because it depends on what kind of equipment the attacker uses.

Finally, some researchers propose to leverage secret keys between the IMD and the readers, using physiological data (such as inter-pulse timing, heart beat frequency, etc.) [28, 4]. However, the security of these schemes is hard to prove and sometimes dubious. In fact, it has been shown that some of these data can be remotely measured with sensitive probes [14]. Furthermore, it is unclear how much entropy there is in the timing of heart beats and even whether some of the secret cannot be retrieved from other channels.

The solution proposed in our paper can enforce any access distance with a high degree of accuracy and is therefore inherently more secure than previously described proximity-based approaches.

Distance-bounding

Independent of the IMD context, related work has been done in the field of location verification based on radio and ultrasonic distance bounding protocols. The first distance-bounding protocol was described in [2]. Other extensions and optimizations followed in [26, 25, 13, 20]. The use of distance bounding based on ultrasonic communication channel was proposed in [24]. The application of distance bounding to key establishment device pairing was proposed in [27].

Our proposal is similar to [27] in that it equally uses ultrasonic distance bounding for proximity-based message authentication. However our proposal includes a number of design choices specific to the IMD context and considers side-channel attacks that were, so far, not considered in the design and implementation of distance-bounding protocols.

7. CONCLUSION

In this work we proposed a novel proximity-based access control scheme for implantable medical devices (IMDs). Our scheme is based on ultrasonic distance-bounding and enables an implanted medical device to grant access to its resources only to those devices that are in its close proximity. It is based on a new proximity-based message authentication protocol based on ultrasonic distance bounding that can authenticate arbitrary messages. We demonstrated the feasibility of our approach through tests in an emulated patient environment. We showed that, although implanted, IMDs can successfully verify the proximity of other devices with a high degree of accuracy. We show that our protocol can be seamlessly combined with a number of existing IMD countermeasures including those that prevent battery draining Denial-of-Service attacks. Our proposal equally includes a number of design choices for distance bounding protocols that are specific to the IMD context and considers side channel attacks that were not so far considered in the design and implementation of distance bounding protocols.

8. REFERENCES

[1] USPTO Patent Application 20080044014. Secure telemetric link. `http://www.freshpatents.com/Secure-telemetric-link-dt20080221ptan200800%44014.php?type=description`.

[2] S. Brands and D. Chaum. Distance-bounding protocols. In *Workshop on the theory and application of cryptographic techniques on Advances in cryptology*, pages 344–359. Springer-Verlag New York, Inc., 1994.

[3] Srdjan Capkun and Jean-Pierre Hubaux. Secure positioning of wireless devices with application to sensor networks. In *IEEE INFOCOM*, 2005.

[4] S. Cherukuri, K.K. Venkatasubramanian, and S.K.S. Gupta. Biosec: a biometric based approach for securing communication in wireless networks of biosensors implanted in the human body. In *International Conference on Parallel Processing Workshops*, pages 432–439, Oct. 2003.

[5] Jolyon Clulow, Gerhard P. Hancke, Markus G. Kuhn, and Tyler Moore. So near and yet so far: Distance-bounding attacks in wireless networks. In *European Workshop on Security and Privacy in Ad-Hoc and Sensor Networks (ESAS)*, pages 20–21, September 2006.

[6] Tamara Denning, Kevin Fu, and Tadayoshi Kohno. Absence makes the heart grow fonder: new directions for implantable medical device security. In *HOTSEC'08: Proceedings of the 3rd conference on Hot topics in security*, pages 1–7, Berkeley, CA, USA, 2008. USENIX Association.

[7] Will Dunham. Mp3 player headphones may hinder pacemakers: study. `http://www.reuters.com/article/domesticNews/idUSTRE4A81SS20081109`, November 2008.

[8] K. Fotopoulou and B. W. Flynn. Optimum antenna coil structure for inductive powering of passive RFID tags. In *RFID, 2007. IEEE International Conference on*, pages 71–77, 2007.

[9] E. Freudenthal, R. Spring, and L. Estevez. Practical techniques for limiting disclosure of rf-equipped medical devices. In *Engineering in Medicine and Biology Workshop, 2007 IEEE Dallas*, 2007.

[10] S.K.S. Gupta, T. Mukherjee, and K. Venkatasubramanian. Criticality aware access control model for pervasive applications. In *Pervasive Computing and Communications, 2006. PerCom 2006. Fourth Annual IEEE International Conference on*, pages 5 pp.–257, March 2006.

[11] Daniel Halperin, Thomas S. Heydt-Benjamin, Kevin Fu, Tadayoshi Kohno, and William H. Maisel. Security and privacy for implantable medical devices. *IEEE Pervasive Computing*, 7(1):30–39, 2008.

[12] Daniel Halperin, Thomas S. Heydt-Benjamin, Benjamin Ransford, Shane S. Clark, Benessa Defend, Will Morgan, Kevin Fu, Tadayoshi Kohno, and William H. Maisel. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. *Security and Privacy, IEEE Symposium on*, 0:129–142, 2008.

[13] G. P. Hancke and M. G. Kuhn. An RFID distance bounding protocol. In *IEEE SecureComm*, 2005.

[14] C J Harland, T D Clark, and R J Prance. Electric potential probes - new directions in the remote sensing of the human body. In *Measurement Science and Technology*, volume 13, page p 163, 2002.

[15] C. Israel and S. Barold. Pacemaker systems as implantable cardiac rhythm monitors. In *American Journal of Cardiology*, pages pp. 442–445, 2001.

[16] Tim Kindberg and Kan Zhang. Validating and securing spontaneous associations between wireless devices. In *ISC*, pages 44–53, 2003.

[17] News-Medical.net. Implantable medical devices may expose patients to security, privacy risks; solutions suggested. `http://www.news-medical.net/?id=36219`.

[18] B. Park, A. D. Whittaker, R. K. Miller, and D. S. Hale. Predicting intramuscular fat in beef longissimus muscle from speed of sound. In *Journal of Animal Science*, pages 109–116. ACM, 1994.

[19] Nissanka Bodhi Priyantha, Anit Chakraborty, and Hari Balakrishnan. The Cricket Location-Support System. In *6th ACM MOBICOM*, Boston, MA, August 2000.

[20] Kasper Bonne Rasmussen and Srdjan Čapkun. Location privacy of distance bounding protocols. In *CCS '08: Proceedings of the 15th ACM conference on Computer and communications security*, pages 149–160, New York, NY, USA, 2008. ACM.

[21] Melanie R. Rieback, Bruno Crispo, and Andrew S. Tanenbaum. RFID guardian: A battery-powered mobile device for RFID privacy management. In *Proc. 10th*

Australasian Conf. on Information Security and Privacy (ACISP 2005), volume 3574 of *LNCS*, pages 184–194. Springer-Verlag, July 2005.

[22] P. Inchingolo S. Bergamasco, M. Bon. Medical data protection with a new generation of hardware authentication tokens. In *Mediterranean Conference on Medical and Biological Engineering and Computing*, 2001.

[23] Bill Saltzstein. Bluetooth wireless technology in the medical market. In *Bluetooth Developers Conference*, December 2001.

[24] N. Sastry, U. Shankar, and D. Wagner. Secure Verification of Location claims. In *Proceedings of the ACM Workshop on Wireless Security (WiSe)*, pages 1–10. ACM Press, September 2003.

[25] Dave Singelée and Bart Preneel. Distance bounding in noisy environments. In *ESAS*, pages 101–115, 2007.

[26] S. Čapkun, L. Buttyán, and J.-P. Hubaux. SECTOR: Secure Tracking of Node Encounters in Multi-hop Wireless Networks. In *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, Washington, USA, October 2003.

[27] Srdjan Čapkun and Mario Čagalj. Integrity regions: authentication through presence in wireless networks. In *WiSe '06: Proceedings of the 5th ACM workshop on Wireless security*, pages 1–10. ACM, 2006.

[28] K. Venkatasubramanian and S. Gupta. Security for pervasive healthcare. In *Security in Distributed, Grid, Mobile, and Pervasive Computing*, pages pp. 349–366, 2007.