Zero Reconciliation Secret Key Generation for Body-Worn Health Monitoring Devices

Syed Taha Ali
University of New South Wales
Australia
taha@student.unsw.edu.au

Vijay Sivaraman
University of New South Wales
Australia
vijay@unsw.edu.au

Diethelm Ostry
ICT Centre, CSIRO
Australia
diet.ostry@csiro.au

ABSTRACT

Wearable wireless sensor devices are key components in the emerging technology of personalized healthcare monitoring. Medical data collected by these devices must be secured, especially on the wireless link to the gateway equipment. However, it is difficult to manage the required cryptographic keys, as users may lack the awareness or requisite skills for this task. Alternatively, recent work has shown that two communicating devices can generate secret keys derived directly from symmetrical properties of the wireless channel between them. This channel is also strongly dependent on positioning and movement and cannot be inferred in detail by an eavesdropper. Existing schemes, however, yield keys with mismatching bits at the two ends, requiring reconciliation mechanisms with high implementation and energy costs that are unsuitable for resource-poor body-worn devices.

In this work we propose a secret-key generation mechanism which uses signal strength fluctuations caused by incidental motion of body-worn devices to construct shared keys with near-perfect agreement, thereby avoiding reconciliation costs. Our contributions are: (1) we analyse channel measurement asymmetries caused by non-simultaneous probing of the channel by the link end-points, (2) we propose a practical filtering scheme to minimize these asymmetries, dramatically improving signal correlation between the two ends without reducing entropy, and (3) we develop a method to restrict key generation to periods of channel fluctuation, ensuring near-perfect key agreement. To the best of our knowledge, this work is the first to demonstrate the feasibility of generating high quality secret keys with zero reconciliation cost in body-worn networks for healthcare monitoring.

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection—*physical security, unauthorized access*

General Terms

Experimentation, Performance, Security

WiSec'12, April 16–18, 2012, Tucson, Arizona, USA.
Copyright 2012 ACM 978-1-4503-1265-3/12/04 ...$10.00.

Keywords

Body Area Networks, Secret Key Generation

1. INTRODUCTION

Soaring national health expenditures and escalating age-related disabilities are shifting the emphasis from the hospital to the home. Body area networks are at the forefront of emerging technologies in this trend towards personalized healthcare. A body area network typically consists of small sensors mounted on the body to record vital signs and communicate wirelessly with a base-station (a fixed access point or a portable device such as a mobile phone) for real-time analysis and possibly remote diagnosis. Wearable platforms for health monitoring have begun to appear in the market. Apple has recently patented a sensor strip device [1] that interfaces with the iPhone, and IMEC has demonstrated a sensor device [2] which communicates with phones running the Android OS. Fig. 1 illustrates a topology based on the Sensium Digital Plaster [3], a body-worn wireless solution to monitor a subject's ECG, temperature, blood glucose and oxygen levels. A report [4] by ABI forecasts that the market for wearable wireless sensor devices will grow to more than 420 million devices by 2014. Securing these devices is a significant challenge considering their low power and computation capabilities, but it is also critical, since these devices record and handle medical data which comes with stringent privacy and liability concerns.

The high computational cost of asymmetric cryptography precludes its use in the body-worn device for encrypting medical data, leaving symmetric encryption using a shared key as the only viable option. The challenge lies in dynamically sharing a secret key between the body-worn device and the base-station. The secret key cannot be pre-configured at time of manufacture, since the pairing of body-worn device to base-station is done at deployment, and dynamic pairing requires a trusted third-party to store the keys, carrying with it risk of compromise and associated liability. Furthermore, experience has shown [5] that users (such as the elderly) are often unaware of the need, or unable to configure secrets of sufficient strength, or safeguard these secrets adequately. It is far more practical to automatically generate secret keys as needed. Moreover, keys need to be renewed periodically to protect against attack. It is more straightforward to generate shared secret keys using the Diffie-Hellman key exchange but it is expensive to implement and execute on resource-constrained sensor devices [6].

Recent work such as [7, 8] has shown that it is possible to generate a shared secret over an unsecured wireless channel

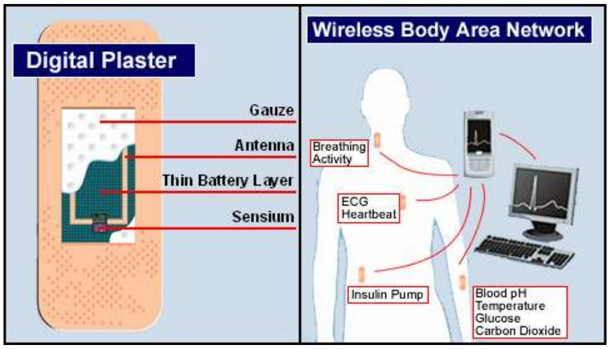

Figure 1: Toumaz Sensium^{TM} Digital Plaster and body area network topology

by exploiting the directional symmetry of the wireless link. Specifically, the multipath propagation characteristics between two communicating parties, Alice and Bob, are symmetric (and hence strongly correlated) at both ends of the link, and yet sufficiently random to allow Alice and Bob to generate shared secret bits. The focus of much of the prior work has been to generate secret bits at a high rate (tens of bits per second), which comes at the cost of more frequent channel probing and greater bit mismatch between the two ends. Even a 2% probability of bit mismatch means that a 128-bit key has only a 7.5% chance of matching perfectly. To resolve mismatch, reconciliation methods such as Cascade [9] are proposed, where the two ends exchange messages to probabilistically identify mismatching bits.

In contrast, we focus on low-data-rate patient monitoring applications that require periodic key renewal. Pairwise temporal keys (or session keys) are advocated in the emerging IEEE 802.15 standard [10] for body area networks. In these applications a high bit generation rate is not essential; for example, if a 128-bit key needs to be renewed every hour, a generation rate of a few bits per minute suffices. This low bit-rate requirement has three benefits for low-complexity key generation schemes. First, the mismatch of key bits generated by the two ends can be avoided to eliminate reconciliation overheads which consume precious computing and communication resources [11]. Second, body-worn devices typically embed their logic in hardware as a single-chip solution (as in the case of the Sensium [3]), and interactive reconciliation protocols requiring real-time communication are too complex to be completely implemented in custom hardware and their flexibility is limited. A third advantage is that the low bit-rate requirement allows the key generation mechanism to piggyback channel sampling on regular data exchanges (typically at rates of the order of 1 packet/s), instead of requiring dedicated channel sounding messages. This significantly reduces radio usage, usually the most expensive operation in small sensor devices.

In this paper, we undertake an experimental study of secret key generation in the specific setting of body-worn devices, and propose a cost-effective scheme to eliminate mismatch between the two ends. Our target is to have at least a 75% chance of generating a fully matching 128-bit secret key, corresponding to bit-agreement probability of at least 99.8%. Our specific contributions are:

1. Our first contribution is a demonstration that the dominant cause of the observed channel mismatch between the two ends of the link during motion is the time delay between measurements by the two ends. We present a theoretical bound on the mismatch, and validate it via experiments with body-worn devices in a representative office environment as well as an anechoic chamber.

2. Our second contribution is a method to reduce this mismatch by filtering the signal using a practical, low-complexity approach that dramatically improves correlation between the two endpoints, without reducing signal randomness.

3. Our third contribution is a mechanism to confine bit generation to periods of high motion-related fluctuation, further reducing disagreement in channel estimation thereby virtually eliminating key-bit mismatch. We show that an activity threshold can be adjusted to yield near-perfect key agreement by trading-off against key generation rate.

For our threat model, we situate passive eavesdroppers at various points in the environment who sample the channel at the same time as the communicating parties and know the key extraction algorithm and its settings. We do not address the issue of authentication in this paper: we believe that establishing initial trust between two parties is a distinct research problem and it is important during the bootstrapping phase, whereas our focus is on key renewal. If we assume a mechanism for bootstrapping initial trust, a basic challenge-response protocol can ensure authenticity of newly generated session keys.

We believe our work is the first to undertake secret-key generation using the wireless channel in the important and unique context of body-worn healthcare devices. Moreover, our scheme dispenses with reconciliation and dedicated channel sampling, while generating high entropy secret bits at a usable rate of approximately 8 bits/min, with 99.8% agreement. At this rate, a usable 128-bit key is generated every 20 minutes. If a session key is renewed over a greater period, say 1 hour, as is recommended for WiFi [12], the probability of generating a perfectly matching key at both endpoints using our mechanism can be up to 99.5%. Our scheme is light-weight, implementable on the current generation of body-wearable devices, and suitable for large-scale deployment in home-based personalized healthcare systems.

The rest of this paper is organized as follows: Section 2 discusses prior work and the key reconciliation process. In Section 3 we identify the cause of mismatch theoretically and experimentally, and in Section 4 describe a filtering technique to minimize it. Section 5 details our region selection and key generation mechanisms, whose performance is then analysed in Section 6. We conclude in Section 7.

2. BACKGROUND

In this section we briefly describe secret key generation and prior contributions, and highlight how our work differs in that it eliminates key reconciliation.

2.1 Secret Key Generation

2.1.1 The Basic Principle

The wireless channel is intrinsically symmetrical by the reciprocity property of electromagnetic propagation. In the absence of interference, noise, and changes in the channel, two communicating parties, Alice and Bob, using identical transceivers and antennas, and transmitting identical signals, will both also receive identical signals. In the complex geometry typical of interior environments, radio signals can propagate via multiple paths, each experiencing a different delay, attenuation, and phase and polarisation distortions which depend on the details of each path. The set of parameters defining the effects of all these paths can be measured by both Alice and Bob and ideally they will agree.

In the time domain, the channel can be represented by the delay spectrum or impulse response, and equivalently by the frequency spectrum in the frequency domain. Measurements of either of these representations can be used by Alice and Bob to construct a shared key, unique to their positions. An eavesdropper, Eve, located outside a distance greater than about one radio wavelength from either Alice or Bob, will measure a different spectrum, and so will be unable to determine their key. This scenario leads to the well-known Jake's uniform scattering model [13] which states that there is rapid decorrelation in the signal over a distance of approximately half a wavelength, and one may assume independent signals for a separation of one to two wavelengths or more.

Measurement of either delay or channel spectra with sufficient resolution to generate long keys requires significant investment in hardware and energy consumption. An approach more suited to energy-constrained devices uses a time series of received signal strengths measured along a trajectory traversed by one or both parties as a source of shared information [7, 8].

In practice, asymmetric components appear in these channel measurements due to transceiver differences, random noise, changes caused by motion, either of the parties or other elements of the environment, and asymmetrically located interference sources. These asymmetries cause discrepancies in the derived keys, requiring additional operations to obtain key agreement.

2.1.2 The Procedure

The process of shared secret key generation described in the literature typically comprises four phases:

1. *Channel sensing:* Alice and Bob each measure some characteristic of the channel. A time series of received signal strengths during node motion is commonly used [7, 8, 14], although other suitable channel characteristics have also been studied [15, 16, 17].
2. *Quantization:* The measurements are converted into a string of key bits. Approaches based on signal extrema [7, 8] and ranking [18] have been described in prior work.
3. *Reconciliation:* Key bit discrepancies at the two ends are discarded or corrected by employing an information reconciliation protocol [19].
4. *Privacy amplification:* The now matching keys are then strengthened by discarding agreed bits or by performing a transformation to increase key entropy and obfuscate any partial information an eavesdropper may have gathered during key reconciliation.

2.1.3 Reconciliation

We now consider the reconciliation phase with a view to showing that it incurs an unjustifiably high cost in body-worn devices, thereby motivating the study in this paper. Information reconciliation mechanisms have been developed mainly in the context of quantum cryptography [19], and key generation schemes for wireless links either borrow these mechanisms or propose non-optimal ad hoc schemes.

To reconcile bitstrings, two parties exchange metadata, (similar in concept to the cyclic redundancy check (CRC)), to identify mismatching bits, whilst simultaneously trying to minimize the potential leakage of information about the bitstring to an eavesdropper. Once mismatching bits are identified, they are either discarded from the bitstring, or else corrected, which may require further message exchanges. Unfortunately, like CRC, reconciliation methods only detect and correct a specific class of errors, with a probability depending on the capabilities of the reconciliation mechanism. If we consider a simple reconciliation scheme which computes a single parity bit over a block, an even number of errors will go undetected. Considering a block of b bits, let q denote the probability that an individual bit differs at both ends. The probability P_q of having mismatching blocks in spite of reconciliation can be expressed as:

$$P_q = \sum_{i=1}^{\lfloor b/2 \rfloor} \binom{b}{2i} (1-q)^{b-2i} q^{2i}$$

Consequently, the probability P, of agreeing on an error-free key of length K is

$$P = (1 - P_q)^{K/b}$$

For example, if there is as little as a 2% chance of a bit mismatching between endpoints, for a block size $b = 8$ there is approximately a 15% chance of uncorrected errors in a key of length $K = 128$, in which case the key will have to be regenerated. And typically, to counter the information leaked to an adversary due to parity bits being exposed, an equal number of bits needs to be dropped from the key, thereby reducing the final key bit rate.

Reconciliation protocols such as Cascade typically perform this parity check multiple times and shuffle the bit sequence in coordination before each test. This incurs significant memory and transmission overheads (as documented in [11]), much more pronounced in the context of a miniature sensor device operating with constrained resources. Furthermore, reconciliation will also add to design complexity, of particular concern since these protocols will typically be implemented in ASICs to provide a single-chip solution for body-worn devices. Our aim, therefore, is to virtually eliminate the need for reconciliation by aiming for a bit agreement ratio of 99.8% or greater, so that a typical 128-bit key has a very good ($> 75\%$) chance of matching perfectly.

2.1.4 Performance Metrics

The following metrics are commonly used to evaluate the performance of secret key generation schemes:

1. *Key Agreement:* the fraction of bits matching at both ends, ideally 100%. Eavesdroppers should match in only about 50% of the bits they generate.
2. *Secret Bit Rate:* the average number of secret key bits extracted from the channel per unit time. This depends on factors such as sampling rate, quantizer parameters, and channel variability.
3. *Entropy:* a measure of the uncertainty (inherent randomness) in the key. A typical measure of entropy of a random variable X, over the set of n symbols $x_1, x_2, ..., x_n$, is

$$H(X) = -\sum_{i=1}^{n} p(x_i) \log_2 p(x_i)$$

where $p(x_i)$ is the probability of occurrence of symbol x_i.

For binary symbols, a value close to 1 indicates high entropy. We use the NIST test suite [20] to estimate entropy.

Ideally a scheme should generate keys with high agreement, at a fast rate, and with high entropy. However, these are conflicting goals and researchers generally focus on one and employ secondary means to improve the others, at additional computational and communication cost. Sampling at a high rate will yield a higher bit rate, but will have greater disagreement, and lower entropy, since the signal variation is lower relative to the sampling rate so that successive bits will be more correlated. Sampling at larger intervals improves key agreement and entropy but reduces bit rate. These tradeoffs are handled in a variety of ways in prior work.

2.2 Prior Work

Prior work in secret key generation for **802.11 WiFi** considers both static and mobile cases. The authors of [7] show that with modified 802.11 hardware able to measure channel impulse response it is possible to obtain keys at a rate of more than 1 bit/s with almost perfect agreement, but use of simple signal strength measurements instead resulted in key disagreements. [8] presents experimental results for several static and mobile scenarios including walking and bicycle-riding. Motion is seen to yield high entropy keys at a high rate and with good key agreement. The authors' emphasis is on high bit generation rates and relatively high bit mismatch is seen (4-30%) making a reconciliation mechanism (Cascade [19]) necessary, along with privacy amplification.

In [21], the authors consider key generation in **ultra wideband channels**, mainly using simulations of static deployments. They use the envelope of the observed channel impulse response, rather than the received signal strength metric. However, successive key values were highly correlated and they use a whitening process employing training data for privacy amplification.

Wireless sensor devices have been specifically considered in some prior work. In [22], the authors measure at a sequence of frequencies to estimate the spectrum and extract keys with agreement of over 97% in static deployments.

In [18], the authors aim for a very high rate key generation of 22 bits per second with 2.2% disagreement, or, alternately 3 bits per second with 0.04% disagreement. The channel is sampled at a rate of 50 probes/s. Extensive processing is done on the data, including interpolation, de-correlation, and multi-bit adaptive quantization. One of the endpoints must be moved continuously in a 'random' manner to induce signal fading fluctuations. This approach is extended in [14] by introducing a ranking mechanism to remove those asymmetries in the received signal strength indicator (RSSI) traces due to differences in hardware characteristics. Experiments with TelosB motes show a key generation rate of 40 bits/s with 4% disagreement.

Body area networks have been considered in only one work in the literature. The authors in [23] simulate a near-body channel to derive an upper bound of 4 bits/s due to inherent limitations on channel entropy but do not describe an actual key generation process.

2.3 Our Focus

Our focus in this paper is on body-worn devices, which have uniquely different constraints and operating conditions. Channel variation is complex and unpredictable in body area networks due to motion, shadowing effects of the human body and multipath propagation [24]. Moreover, due to the limited resources of body-worn devices, key generation has to be done at minimal cost. In contrast to earlier schemes, we forego the high costs of dedicated sampling, reconciliation and privacy amplification. Our scheme samples the channel in the course of routine transmissions, controls the prime source of bit discrepancies using low-complexity filtering, and relies on the user's own motion to create channel entropy which is harnessed for secret key generation.

3. UNDERSTANDING DISAGREEMENT

In this section, we use theoretical and experimental approaches to show that non-simultaneous sampling of the channel contributes significantly to disagreement between the two ends of the link.

3.1 Theoretical Estimation of Disagreement in Measurements of Link Signal Power

Here we carry out a simplified analysis to estimate the effects of motion on the received signal power measured by the nodes at the ends of a link. There are three well-known contributors to changes in signal power caused by node motion [25]: (i) *path loss*, due to geometric signal spreading has an inverse-square law relationship with range, (ii) *shadow* or *large-scale fading*, arising from signal blockage in the environment including the subject's body and from changes in antenna orientation which affect signal strength through the antenna radiation pattern, and (iii) *small-scale fading*, due to signal fluctuations caused by motion induced changes in the multiple propagation paths between the two nodes. At speeds typical of human motion, range (path loss) and orientation (shadow fading) cause only slow variations in signal strength over successive packets. However multipath (small-scale fading) can cause rapid fluctuations in signal strength as a node's position changes. These components are illustrated in Fig. 2.

Consider an environment with appreciable multipath propagation, i.e. where multiple propagation paths exist between the two nodes: suppose at time instant $t = 0$, the stationary node (the base-station (BS)) samples the channel (i.e. hears a transmission from the mobile node), and Δt seconds later the (body-worn) mobile node samples the channel (i.e. hears the transmission from the BS). The difference in channel measurements between the two end-points is equivalent to the change in channel from time $t = 0$ to time Δt as measured by one node (say the mobile node) at the two instants, since the channel is reciprocal at each time. In what follows we estimate this change using a simple model.

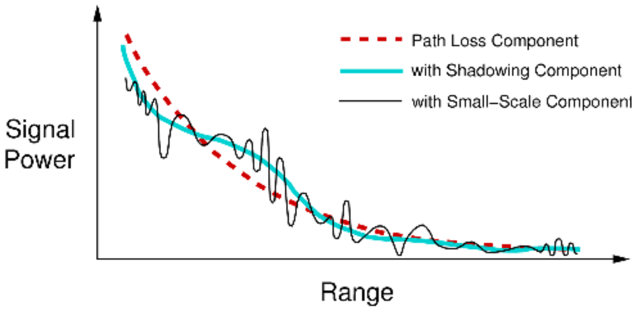

Figure 2: Components of the received signal power

When the BS transmits, signals propagating along the multiple paths combine to form a standing wave pattern in the environment. At places where the signals reinforce due to phase agreement, there is an increase of signal strength, and at places where the signals subtract there is a decrease in signal strength. As the mobile node moves through the environment, the signal strength it observes fluctuates due to these interference effects. Because of the fixed characteristic radio signal wavelength, adjacent locations where the signal is maximum or minimum cannot be separated by less than a distance of the order of half a wavelength [26]. This places an upper bound on the rate at which the received signal power can change as the node moves through the standing wave pattern. If the signal radio wavelength is λ and the receiver moves at velocity v, the maximum frequency at which the observed signal power can change in the receiver is

$$f_{\max} = v \cdot (2/\lambda). \tag{1}$$

This bound limits the worst-case (i.e. highest frequency) signal component that the receiver senses to

$$y(t) = (A/2)\sin 2\pi f_{\max} t. \tag{2}$$

where A is the peak-to-peak amplitude of the signal. The maximum discrepancy in amplitude, Δy between sample points taken Δt apart in time occurs at $t = 0$ and is

$$\begin{aligned} \Delta y &\approx dy/dt \cdot \Delta t \\ &= (A/2)\cos(2\pi f_{\max} t) 2\pi f_{\max} \Delta t \\ &= A\pi f_{\max} \Delta t, \quad \text{at } t = 0. \end{aligned} \tag{3}$$

The fractional discrepancy $\epsilon = \Delta y/A$, namely the change as a fraction of the amplitude, is then

$$\begin{aligned} \epsilon &= \pi f_{\max} \Delta t \\ &= 2\pi v \Delta t/\lambda. \end{aligned} \tag{4}$$

At an operating frequency of 2.4GHz for example (where $\lambda = 0.125$m) and a node velocity of $v = 1$m/s, a $\Delta t =$ 20ms delay between the two ends in sampling the channel leads to a maximum fractional error of $\epsilon \approx 1$, implying that the signal component due to changing multipath (excluding contributions due to variation in range and orientation) may have changed over the entire range from a minimum to a maximum during that interval. Since typical wireless sensor network radios today (e.g. the CC2420 [27]) take 20-40ms to measure the wireless link in two directions, this error can be significant in practice and can lead to mismatch between the two ends, as will be examined experimentally next.

3.2 Experiments in Indoor Environment and Anechoic Chamber

We studied two environments experimentally: a representative indoor office environment, and an RF anechoic chamber with very low reflections. The purpose of the experiments is two-fold: (1) to verify the importance of the small-scale fading component (due to multipath) on channel measurement mismatches between the two ends, and (2) to show the effect of channel sampling delay Δt on measurements at the two ends.

Our experiments used MicaZ motes running TinyOS and operating in the 2.4 GHz band. Their radios output a received signal strength indicator (RSSI), a measure of signal power in logarithmic units, related in a simple way to dBm. Our setup is modeled after a real body area network where the body-worn node (Alice), shown in Fig. 3(a)), transmits one packet per second, a rate typical for a health monitoring device sending patient physiological information such as heart-rate, ECG, etc. Even though continuous patient monitoring devices may collect medical readings several times per second, they usually process them in-node (e.g. by av-

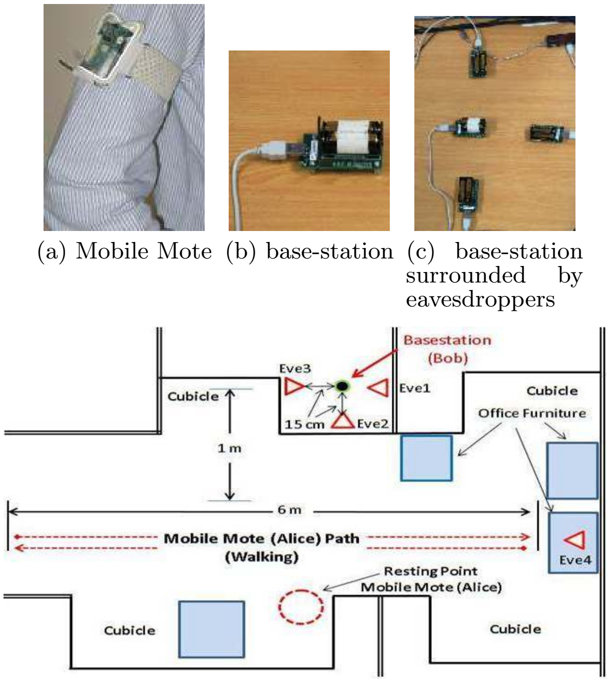
(a) Mobile Mote (b) base-station (c) base-station surrounded by eavesdroppers

(d) Experimental setup of indoor environment

Figure 3: Mobile node, base-station and experimental layout for indoor environment

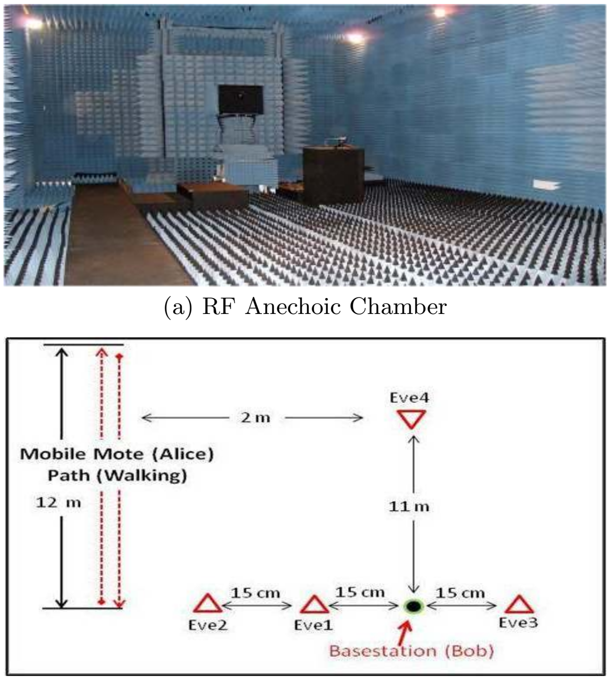
(a) RF Anechoic Chamber

(b) Experimental setup of anechoic chamber

Figure 4: Anechoic chamber and layout

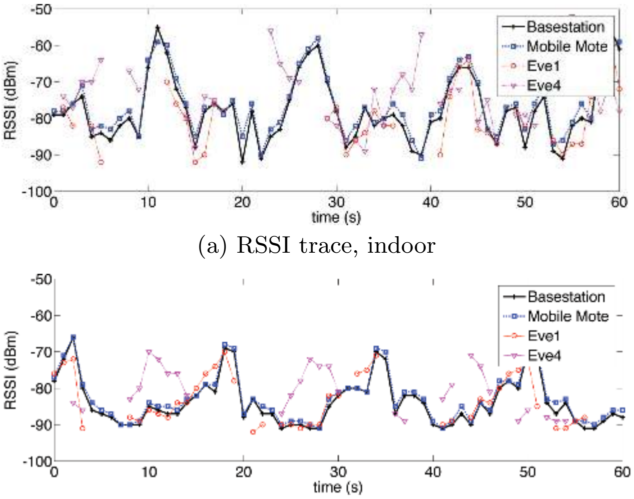

(a) RSSI trace, indoor

(b) RSSI trace, anechoic chamber

Figure 5: Results for Indoor Office and Anechoic Chamber

eraging or aggregating), and then transmit the result to the base station, thereby reducing radio usage. The base-station (Bob, shown in Fig. 3(b)) responds with an acknowledgement as soon as possible (typically 10-20ms on the MicaZ), and this allows the two ends of the link to probe the channel alternately in quick succession.

Our *indoor environment* experiments have the layout depicted in Fig. 3(d) showing the location of the base-station, the four eavesdroppers labeled Eve1 to Eve4, (as shown in Fig. 3(c)), and the path along which the subject walked back and forth. Multiple WiFi networks were operating at the site, but our results did not show evidence of interference. (It is relevant to mention here that efforts are underway to allocate spectrum specifically to body area network applications, to limit interference from other systems [28]).

The RF anechoic chamber is pictured in Fig. 4(a). All surfaces (floors, ceilings, walls) are covered in material that absorbs electromagnetic energy, thereby minimizing RF reflections and consequently the small-scale fading due to multipath propagation. Our experimental layout is shown in Fig. 4(b). In all experiments the subject walked at a moderate pace of about 1m/s.

For the indoor office environment, we show in Fig. 5(a) the signal strengths measured by the base-station, mobile node, and two eavesdroppers (other eavesdroppers show similar results). We observe that the eavesdroppers are not able to replicate the channel measurements accurately, confirming that the base and mobile can use the RSSI measurements to generate random keys. However, we find that there are discrepancies between the signal strengths measured by the base and mobile. The same experimental procedure repeated in the anechoic chamber (which largely eliminates small-scale fading), gave the RSSI trace shown in Fig. 5(b). The signal strength can be seen to vary more smoothly for the base-station and mobile node as compared to the office environment, and correlates better between the two ends.

We examine the discrepancy (i.e. difference in RSSI between the two ends) more closely in Fig. 6 where a box plot depicts the variance observed by both parties. The central mark is the median, the edges of the box denote the 25th and 75th percentiles, the whiskers extend to the most extreme datapoints, and the outliers are plotted individually. For the indoor environment, the discrepancy is seen to vary by as much as 12dB (-6dB to 8dB), able to cause significant mismatch in key bits between the two ends. The mismatch is clearly much lower in the anechoic chamber (no more than about 4dB).

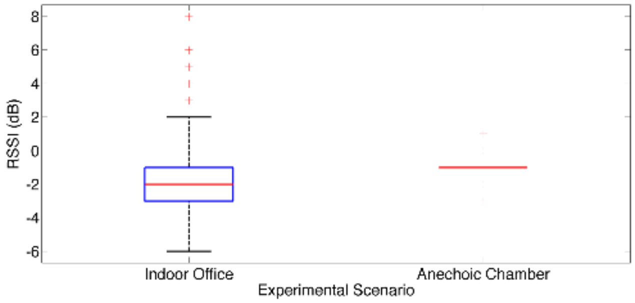

Figure 6: Box plot highlighting discrepancy in RSSI for both test environments

This mismatch can be quantified with the Pearson correlation coefficient r:

$$r = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \bar{X})^2}.\sqrt{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}}$$

where X_i and Y_i are the RSSI values of the ith packet of each party and $\bar{X}$ and $\bar{Y}$ are the respective mean RSSI values of a sequence of n packets. The correlation coefficient r returns a value in $[-1, 1]$ where 1 indicates perfect correlation, 0 indicates no correlation, and -1 indicates anti-correlation. This metric has the benefit that it measures variations and not the absolute values, and so is unaffected by offsets in RSSI measurements arising from different receiver sensitivities or transmit powers. For the indoor office environment, the correlation between the RSSI signals at the base-station and the body-worn node over the entire trace (several minutes) is 0.975, while it is higher, at 0.994, in the anechoic chamber. This provides quantitative confirmation that the multipath (i.e. small-scale fading) component, which occurs in the indoor office environment but is largely absent in the anechoic chamber, is a significant contributor to RSSI mismatch (which leads to key disagreements) between the two communicating parties.

We validate experimentally that mismatch increases with increase in probing delay Δt. We configure the mobile node to acknowledge packet reception from a basestation several times at 40ms intervals. The discrepancy between the RSSI of the original packet (from base to mobile) and the RSSI of each subsequent response (acknowledgement from mobile to base) is measured, and plotted in Fig. 7, for both the indoor office environment and the anechoic chamber. Two observations emerge from this plot: (i) the discrepancy is again much lower in the anechoic chamber than in the indoor office environment, and (ii) the RSSI trace of the first acknowledgement shows least fluctuation, while each subsequent response deviates more (i.e. has larger amplitude). The latter visual observation can be quantified with the correlation coefficient, plotted in Fig. 8 as the probing delay Δt between the two ends increases. It clearly demonstrates that the correlation steadily falls (i.e. mismatch increases) as probing delay increases, and that a 40ms probing delay in the indoor environment is equivalent to a 100ms probing

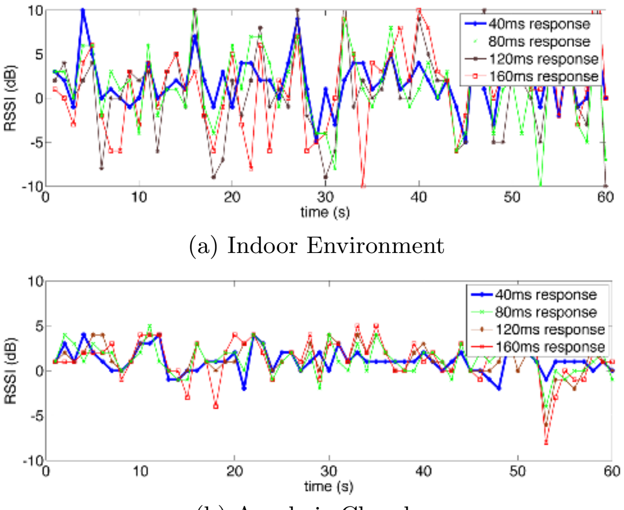

(a) Indoor Environment

(b) Anechoic Chamber

Figure 7: Mismatch due to probing delay Δt

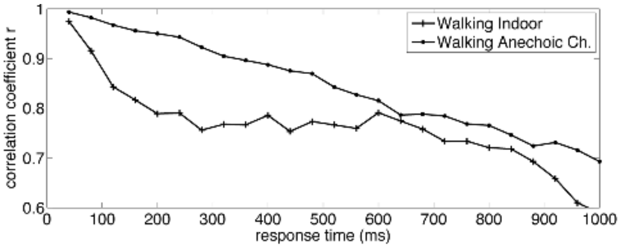

Figure 8: Correlation coefficient r versus sampling delay

delay in the multipath-free anechoic chamber in the sense of yielding a similar correlation of about 0.976.

The theoretical and experimental observations above provide strong evidence that the discrepancy in channel measurement is predominantly due to the lag in sampling by the two ends of the link. In the next section, we develop a novel means of reducing this discrepancy.

We wish to emphasize that other factors such as external interference (which can be asymmetric) and uncorrelated random noise effects (e.g. due to receiver circuitry) also contribute to the discrepancy. To illustrate this, we conducted experiments in which the mobile node is resting at one spot (indicated in Fig. 3(d)), and plot the resulting RSSI in Fig. 9. The channel is relatively static, yet small RSSI discrepancies are visible. Unfortunately these small discrepancies can lead to key mismatch, since the (uncorrelated) noise is amplified by the quantizer to generate key bits. This issue is addressed in Section 5, where we develop a way to eliminate the effects of uncorrelated noise.

4. REDUCING DISAGREEMENT BY FILTERING

In Section 3.1 we developed a simple model showing that the maximum fractional error due to small-scale fading is $\epsilon = \pi f_{max} \Delta t$ where $f_{max} = v \cdot (2/\lambda)$. To reduce this error ϵ, one would ideally like to minimize sampling delay, Δt, but unfortunately the maximum possible reduction is limited by operation in half-duplex mode (although recent proposals for single-channel full-duplex operation [29] may offer a means of overcoming this in future). The other parameter that can be manipulated is the mobile node velocity v, but that would restrict application to slow-moving mobile nodes.

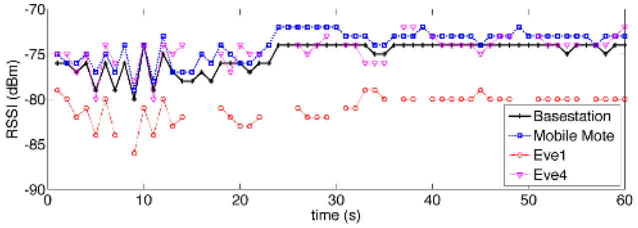

Figure 9: Variation in RSSI for Resting Scenario

Instead, we reduce $f_{\max}$, i.e. the maximum frequency of changes in received signal power arising from motion in a small-scale fading environment. By applying a low-pass filter with cutoff frequency $f_c < f_{\max}$ at both ends of the link, the maximum fractional error in measuring signal power is reduced to $\hat{\epsilon} = \pi f_c \Delta t = \epsilon f_c / f_{\max}$. For the example considered in Section 3.1, where the subject walks at $v = 1$m/s, the delay in bidirectional probing is $\Delta t = 20$ms, and operating frequency is 2.4GHz with wavelength $\lambda = 12.5$cm, we showed that $f_{\max} \approx 16$Hz and the error can theoretically be as high at $\epsilon \approx 100\%$. To restrict this error to less than a desired bound, say $\hat{\epsilon} \approx 3\%$, we can set the filter cut-off frequency to $f_c = (\hat{\epsilon}/\epsilon) f_{\max} \approx 0.48$Hz.

A low-pass Fourier filter is unsuitable for real-life situations where users' motion causes discontinuities and unpredictable changes in the RSSI trace (and is hence not well-modeled by discrete frequency components). Instead we choose the Savitzky-Golay filter [30] which is better able to match the logarithmic form of signal strength measurements given by the receiver RSSI output data. The Savitzky-Golay filter behaves as a low-pass filter [31], and is able to follow the underlying slow-moving features of the RSSI traces we have observed, while providing a controllable reduction in the bandwidth of fluctuations caused by motion in a multipath environment. Moreover, this filter is a linear algorithm that can be easily implemented in ASIC as part of a body-worn solution.

For our experimental work in this paper we select the parameters of the Savitzky-Golay filter for a cut-off frequency $f_c \approx 0.48Hz$, so that the maximum fractional error $\hat{\epsilon}$ is limited to around 3% (as argued above). The mapping of filter parameters to 3 dB cut-off frequency is based on the approximation derived in [31, Eq. (11)]:

$$f_c \approx \frac{K+1}{1.6F - 3.6}, \tag{5}$$

where K is the polynomial order used by the Savitzky-Golay filter, and F is the frame (window) size. We chose $K = 5$ (i.e. 5-th order polynomial) and $F = 11$ (for an impulse response half-length of 5), giving a cut-off frequency $f_c \approx 0.43$Hz, close to the desired value. This filter yielded visually good signals for key generation in all our experiments. Dynamically tuning the filter parameters to adapt to the mobility of the monitored subject is left for future work.

It is important to emphasize that the proposed filtering operation does not reduce the randomness of the signal (and hence of the generated keys). Motion-induced discrepancies occupy a range of frequencies and it is the higher ones, contaminated by the half-duplex delays, which are removed, leaving the lower-frequency components which retain the

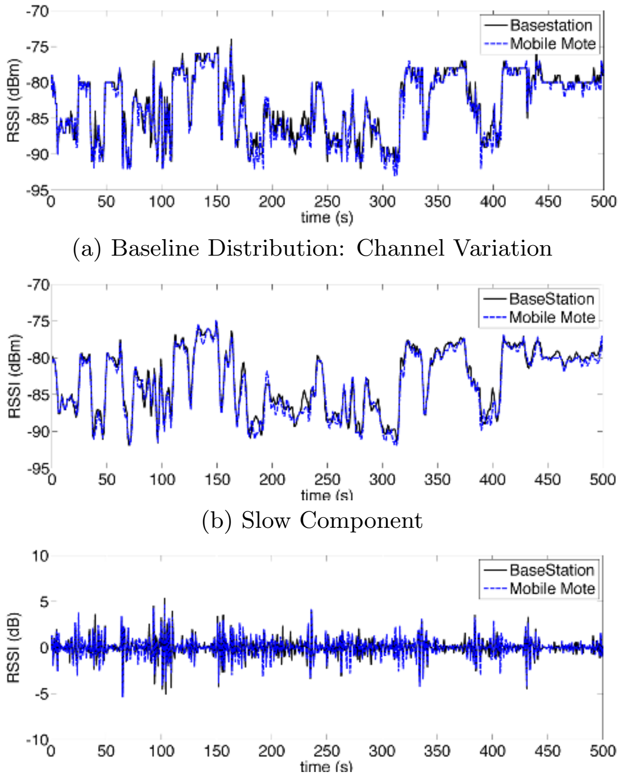

(a) Baseline Distribution: Channel Variation

(b) Slow Component

(c) Fast Component

Figure 10: Application of Savitzky-Golay filter

information about changes in the multipath with position needed for key generation.

To illustrate the operation of the Savitzky-Golay filter, we show its effect in routine subject activity in the indoor office environment over several hours. Fig. 10(a) shows the original RSSI traces. The output of the Savitzky-Golay filter is shown in Fig. 10(b): we call this the *slow component*, and it is primarily attributable to path loss, shadow fading and filtered small-scale fading. The residual (i.e. original signal less the filter output) is shown in Fig. 10(c). We call this the *fast component*, since it consists of higher frequency small-scale fading components which are primarily responsible for the disagreement between the two ends.

Comparing Fig. 10(b) and 10(a), we see that filtering visibly improves agreement between base-station and mobile node. The correlation coefficient of the original RSSI signal between the two ends is 0.973, whereas after filtering, the correlation (of the slow components at the two ends) improves to 0.986. This is almost comparable to the correlation seen in the anechoic chamber, making near-perfect key agreement feasible.

5. DYNAMIC REGION SELECTION AND SECRET KEY GENERATION

We have shown that correlation between the two ends can be greatly improved by filtering the RSSI signals to attenuate the high-frequency components associated with sampling delay. However, factors such as (asymmetric) interference and (uncorrelated) random noise also contribute to mismatch. Indeed the impact of these effects is amplified when the channel is very quiescent (as we showed for a resting subject in Fig. 9), which can lead to an undesirably high rate of secret-key bit-mismatches after quantization. We propose, next, a novel means of dealing with such effects by restricting key-bit generation to periods in which usable signals are available to the quantizer.

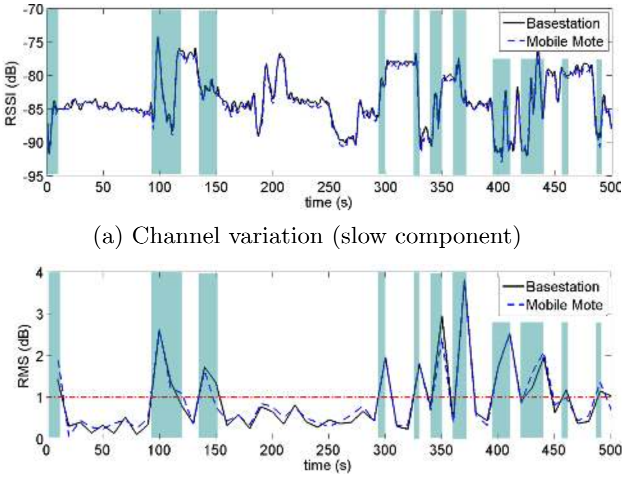

(a) Channel variation (slow component)

(b) Energy of fast component

Figure 11: Region Selection on a trace of routine office activity

5.1 Dynamic Region Selection

When the channel exhibits significant fluctuations (i.e. when the subject is moving rather than resting), the correlated fluctuations in the signal at the two ends have large amplitude and dominate the uncorrelated noise, leading to better agreement (as well as high key entropy). This has been reported in the literature, and indeed some works [14, 8] have explicitly required that the subject should move during key generation. This can place a burden on users, and instead we extend our algorithm to automatically detect time periods (or regions) that are most suitable for secret-key bit generation.

The key observation is that rapid channel variation arises from rapid changes in multipath, and this is strongly expressed in the higher-frequency components of small-scale fading. The latter is already conveniently available to us as the *fast component*, namely the residual between the original and filtered signals. By measuring the RMS energy in the fast component, we can deduce whether there is sufficient activity in the channel for generating high agreement or "good" key-bits at little additional computational cost.

We illustrate this approach in signals obtained in the same office environment as in Section 3 while the subject was engaged in routine office activity. The RSSI (slow component obtained after filtering) is shown in Fig. 11(a), while the RMS energy (in dB, computed using a non-overlapping moving window of $W_{RMS} = 10$ samples) is shown in Fig. 11(b). High energy in the fast component is clearly associated with significant variability in the slow component, and so offers a reliable measure of channel fluctuation. The shaded zones in the figure highlight periods when the fast component energy exceeds a threshold $\theta = 1$dB and dynamically identify regions of high activity during which key bits should be generated from the slow component of the RSSI signal.

5.2 Sampling, Filtering and Quantization

Our threat model considers one or more eavesdroppers (Eve) in the environment who sample the channel at the same time as the legitimate parties, and know the key extraction algorithm and parameters. However, we stipulate that Eve is separated from the two parties by a distance greater than one radio wavelength ($\sim$ 12.5 cm for the 2.4GHz band), and thereby restricted to measuring a different multipath channel. We do not consider here the issue of initial trust between base-station and mobile node, nor that of active attackers engaged in jamming and packet injection.

The key generation mechanism runs as a background process to normal device operation, and the process flow is depicted in Fig. 12 identifying the input variables required at every stage. For all experiments, we employ a sampling rate of $\tau = 1$ sample/s, allowing channel sampling through routine data transmissions and also reducing correlation between successive RSSI readings. The channel response profile is passed to the Savitzky-Golay filter (configured with polynomial order $K = 5$ and frame size $F = 11$ as noted earlier) which outputs the "slow component". The "fast component" is obtained by subtracting the slow component from the original distribution, and its RMS energy is computed for region selection. When periods of high activity (i.e. when the energy exceeds a specified threshold θ) are identified, the corresponding segments of the slow component are passed to the quantizer for bit generation.

Our research does not develop a new quantizer. Instead, we use a basic single-bit quantizer, taken from [7] and refined in [8], and operating as follows: the base-station and mobile node define an adaptive moving window of size W_Q, within which they process blocks of consecutive (filtered) RSSI readings. The process is depicted in Fig. 13. For each block, two threshold values are calculated:

$$q+ = \mu + \alpha.\sigma$$
$$q- = \mu - \alpha.\sigma$$

where μ is the mean, σ is the standard deviation, and $\alpha \geq 0$ is an adjustable parameter. If an RSSI reading within a window is greater than $q+$, it is encoded as 1, and if less than $q-$, as 0. The thresholds define an exclusion zone and values falling between them are discarded. Smaller RSSI variations are more likely to disagree at both endpoints and are therefore not considered, in favor of larger excursions. The α parameter allows the operator to adjust quantizer performance to balance between bit generation rate and mismatch. For our purposes, we use a window size of $W_Q = 5$ and $\alpha = 1$, consistent with prior work.

Once both parties generate enough secret bits to form a key, it can be verified using a challenge-response protocol. If the key fails, it is discarded and the process is repeated until keys agree. Results indicate that in typical conditions, this scheme can generate $2 \sim 4$ usable keys per hour.

6. RESULTS AND ANALYSIS

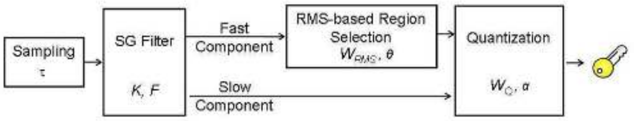

Figure 12: Flow chart of key generation process

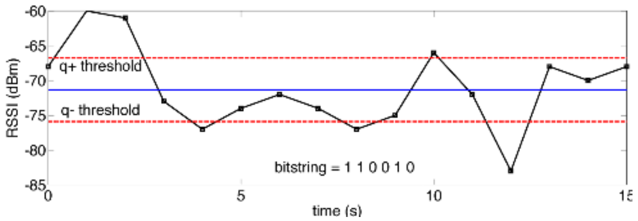

Figure 13: Quantization Process

We tested our key generation mechanism in the office space in Fig. 3(d). The base station is stationary with three eavesdroppers deployed around it at distances of 22cm, 44cm and 100cm. The subject wore the mobile mote on his upper arm. In the first experiment, the subject performed *High Activity*, working, walking and interacting with other people in the room. In the second experiment he performed *Low Activity*, mainly seated at his cubicle working and occasionally getting up to fetch items from other cubicles. Care was taken to ensure the experiments were performed in a realistic manner, as close as possible to an actual deployment of bodyworn sensor devices. Trace data is collected from each experiment for 40 minutes and our key generation scheme is applied offline to assess its performance with different parameter settings.

Table 1 shows, for the *High Activity* scenario, the percentage of key bits that agree for different energy threshold settings. Filtering improved signal correlation between the two ends, but after quantization the key agreement improved only marginally (from 97.27% to 97.91%). This is because the quantizer amplifies uncorrelated random noise during quiescent periods (as explained in Section 4). However, when dynamic region selection is applied to restrict key-bit generation to regions with at least θ dB energy in the fast component, key agreement improves dramatically: a threshold setting of $\theta = 0.5$ dB improves key agreement to over 99%, and at $\theta = 1.5$ dB key bits were found to match with probability over 99.8%.

For *Low Activity*, Table 2 shows that agreement of keys generated from the raw signal is quite low at around 93%. This can be attributed to longer quiescent or low-motion periods during this experiment where the subject just sits at his desk, and uncorrelated noise effects dominate channel variation. Filtering improves this only marginally, but when combined with region selection there is again a dramatic impact. Threshold $\theta = 0.5$ dB improves key agreement to over 98%, while at $\theta = 1.5$ dB key bits were found to match with probability over 99.8%. This demonstrates that filtering and region selection together can effectively improve bit agreement to near-ideal levels.

The high key entropy seen in all cases (> 0.99) (Column 6), and the keys' passing the NIST *approximate entropy* test [20] confirm that the Savitzky-Golay filter retains a sufficient component of the essential randomness arising from motion in a multipath environment.

Filtering and region selection improve bit agreement, but reduce the bit generation rate (Column 2 of the Tables). The Savitzky-Golay filter smooths out the more rapid variations in the raw signal, effectively reducing the number of larger excursions that the quantizer directly maps to key bits. This causes bit generation rate to decrease from 0.33 to 0.24 bits/s (for high activity) and from 0.21 to 0.19 bits/s for low activity. Region selection further reduces bit generation rate, because with increasing threshold, a progressively

Signal quantized	Key Agreement (%)	bit rate (bit/s)	Eve1 Key Agreement (%)	Eve2 Key Agreement (%)	Eve3 Key Agreement (%)	Entropy
unfiltered	97.27	0.33	47.33	47.07	50.42	0.9979
filtered	97.91	0.2435	51.31	51.43	51.20	0.9990
filtered, $\theta = 0.5$	99.08	0.2221	51.68	51.22	51.17	0.9990
filtered, $\theta = 1$	99.74	0.1812	52.03	51.04	51.63	0.9992
filtered, $\theta = 1.5$	99.83	0.1410	51.31	50.66	50.68	0.9992
filtered, $\theta = 2$	99.88	0.1010	50.72	50.74	50.90	0.9995
filtered, $\theta = 2.5$	99.92	0.0647	51.67	52.12	51.27	0.9996
filtered, $\theta = 3$	100	0.0370	51.31	51.18	51.30	0.9997

Table 1: Effect of varying threshold θ on key generation performance metrics for High Activity scenario

Signal quantized	Key Agreement (%)	bit rate (bit/s)	Eve1 Key Agreement (%)	Eve2 Key Agreement (%)	Eve3 Key Agreement (%)	Entropy
unfiltered	93.04	0.2174	49.43	48.89	49.62	0.9970
filtered	93.06	0.1968	49.10	49.04	49.30	0.9993
filtered, $\theta = 0.5$	98.41	0.1323	49.21	48.98	49.43	0.9994
filtered, $\theta = 1$	99.41	0.0861	48.96	48.77	48.24	0.9993
filtered, $\theta = 1.5$	99.80	0.0570	49.51	49.03	50.32	0.9994
filtered, $\theta = 2$	100	0.0365	47.90	49.14	48.54	0.9995

Table 2: Effect of varying threshold θ on key generation performance metrics for Low Activity scenario

smaller proportion of the signal (of high channel activity) is available for quantization. This trade-off is illustrated in Fig. 14, which shows that with increasing threshold θ, the key agreement (left axis) increases while the bit generation rate decreases (right axis), for both high and low activity. A compromise can be chosen by choosing θ appropriately. For the scenarios we considered, a threshold value of $\theta = 1.5$ is sufficient for 99.8% bit agreement, corresponding to a 75% chance of both endpoints' agreeing on a 128-bit secret key. For this threshold setting, our scheme achieves a bit rate of $0.057 \sim 0.141$ bits/s, i.e. it would take $15 \sim 35$ minutes to generate a usable 128-bit key, quickly enough for key renewal purposes. If typical session key lifetime is approximately 1 hour, the chances of having a valid new key perfectly matching at both endpoints varies from $93.5 \sim 99.5\%$ depending on the user's activity.

Tables 1 and 2 also show the percentage of matching bits that each of the eavesdroppers generate by passively listening to the channel. Eavesdropper agreement hovers near 50% for all cases, which is ideal, indicating that their chance of guessing if a generated bit is correct or not is the same as an unbiased coin toss.

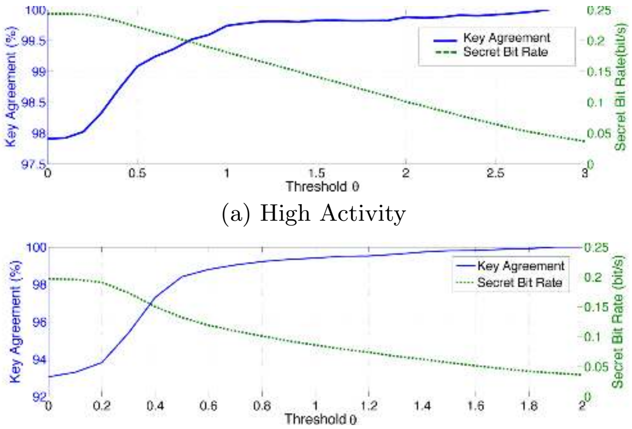

(a) High Activity

(b) Low Activity

Figure 14: Key agreement vs. secret bit rate for varying region selection threshold θ

7. CONCLUSION

In this paper we presented a method for generating shared secret keys using motion in body area networks. Our first contribution has been to identify the cause of key mismatch: we presented a theoretical model to account for the mismatch in secret-key agreement, and validated it with experiments in an indoor environment and in an anechoic chamber. Our results showed that disagreement between two endpoints is primarily due to half-duplex measurement delays and random noise. Furthermore, we noted that these discrepancies are concentrated in the rapidly-varying component of the channel RSSI trace. Second, we showed that this component can be removed using the Savitzky-Golay filter to dramatically improve endpoint correlation. Our final contribution demonstrated how this residual fast component can be employed to dynamically identify regions of high channel variability, where near-perfect key agreement occurs. Our mechanism is low-cost, does not require dedicated channel sampling or information reconciliation, and incrementally generates high entropy key bits at a rate suitable for key renewal. We used our key generation solution in a real office environment and showed that it takes $15 \sim 35$ minutes to generate a 128 bit key with a 75% chance of perfect agreement between endpoints. If the typical lifetime of a session key is one hour, depending on the subject's activity, there is a $93.5 \sim 99.5\%$ chance of both parties generating a perfectly matching secret-key.

For future work, we intend to study multi-party key agreement for body area networks, and research safeguards against active attackers.

8. ACKNOWLEDGMENTS

The authors would like to thank Mr. Ken Smart of CSIRO and Mr. Linjia Yao for assistance with experiments in the RF anechoic chamber.

9. REFERENCES

[1] Apple Inc. *Sensor Strip.* http://www.patentlyapple.com/patently-apple/2010/03/body-area-networks-apple-sensor-strips-the-iphone.html.

[2] D. Graham-Rowe. Body Organs can Send Status Updates to Your Cellphone. New Scientist, October 2010.

[3] Toumaz Technology Ltd. *Sensium Life Platform.* http://www.toumaz.com/page.php?page=sensium_intro.

[4] ABI Research Service. *Market for Wearable Wireless Sensors to Grow to More than 400 Million Devices by 2014*, 2009. http://www.abiresearch.com.

[5] Bruce Schneier. MySpace Passwords Aren't So Dumb. *WIRED*, December 2006.

[6] E. Blass and M. Zitterbart. Efficient Implementation of Elliptic Curve Cryptography for Wireless Sensor Networks. Technical report, Universität Karlsruhe, 2005.

[7] S. Mathur, W. Trappe, N. Mandayam, C. Ye, and A. Reznik. Radio-telepathy: Extracting a Secret Key from an Unauthenticated Wireless Channel. In *ACM MobiCom*, 2008.

[8] S. Jana, S. N. Premnath, M. Clark, S. Kasera, N. Patwari, and S. Krishnamurthy. On the Effectiveness of Secret Key Extraction Using Wireless Signal Strength in Real Environments. In *ACM MobiCom*, Beijing, 2009.

[9] G. Brassard and L. Salvail. Secret-key Reconciliation by Public Discussion. In *EUROCRYPT*, 1994.

[10] IEEE 802.15 WPAN Task Group 6. *MedWiN MAC and Security Proposal Documentation*, September 2009.

[11] P. Bellot and M. Dang. BB84 Implementation and Computer Reality. In *IEEE RIVF*, 2009.

[12] Tim Moore. IEEE 802.11-01/610r02: 802.1.x and 802.11 Key Interactions. Technical report, Microsoft Research, 2001.

[13] W. C. Jakes. *Microwave Mobile Communications*. Wiley, 1974.

[14] J. Croft, N. Patwari, and S. Kasera. Robust Uncorrelated Bit Extraction Methodologies for Wireless Sensors. In *ACM/IEEE IPSN*, 2010.

[15] B. Azimi-Sadjadi, A. Kiayias, A. Mercado, and B. Yener. Robust Key Generation from Signal Envelopes in Wireless Networks. In *ACM CCS*, 2007.

[16] A. Sayeed and A. Perrig. Secure Wireless Communications: Secret Keys through Multipath. In *IEEE ICASSP*, 2008.

[17] N. Patwari and S. K. Kasera. Temporal Link Signature Measurements for Location Distinction. *IEEE Transactions on Mobile Computing*, 10(3):449–462, March 2011.

[18] N. Patwari, J. Croft, S. Jana, and S. K. Kasera. High Rate Uncorrelated Bit Extraction for Shared Key Generation from Channel Measurements. *IEEE Transactions on Mobile Computing*, 9(1), 2010.

[19] G. Brassard and L. Salvail. Secret-Key Reconciliation by Public Discussion. In *Workshop on the Theory and Application of Cryptographic Techniques (EUROCRYPT)*, 1994.

[20] NIST. *A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications*, 2001.

[21] R. Wilson, D. Tse, and R. A. Scholtz. Channel Identification: Secret Sharing using Reciprocity in Ultrawideband Channels. *IEEE Transactions on Information Forensics and Security*, 2(3), 2007.

[22] M. Wilhelm, I. Martinovic, and J. B. Schmitt. Secret Keys from Entangled Sensor Motes: Implementation and Analysis. In *ACM WiSec*, 2010.

[23] L. W. Hanlen, D. Smith, J. Zhang, and D. Lewis. Key-sharing via Channel Randomness in Narrowband Body Area Networks: Is Everyday Movement Sufficient? In *Bodynets*, 2009.

[24] David Smith, Leif Hanlen, Andrew Zhang, Dino Miniutti, David Rodda, and Ben Gilbert. First and Second-Order Statistical Characterizations of the Dynamic Body-Area Propagation Channel of Various Bandwidths. *Annals of Telecommunications*, 66(3-4):187–203, 2011.

[25] B. Sklar. Rayleigh Fading Channels in Mobile Digital Communication Systems. *IEEE Communications Magazine*, 35(7), 1997.

[26] R.P. Bowman. Quantifying Hazardous Microwave Fields. In *Microwave Bioeffects and Radiation Safety.* University of Alberta, Canada: International Microwave Power Institute, 1978.

[27] ChipCon Products. *2.4 GHz IEEE 802.15.4 / Zigbee-ready RF Transceiver.*

[28] G. Lawton. More Spectrum Sought for Body Sensor Networks. *Computing Now*, Oct. 2009.

[29] J. I. Choi, K. Srinivasan, M. Jain, P. Levis, and S. Katti. Achieving Single Channel, Full Duplex Wireless Communication. In *ACM MobiCom*, 2010.

[30] A. Savitzky and M. J. E. Golay. Smoothing and Differentiation of Data by Simplified Least Squares Procedures. *Analytical Chemistry*, 36:8, 1964.

[31] R. W. Schafer. On the Frequency-Domain Properties of Savitzky-Golay Filters. Technical report, HP Laboratories, HPL-2010-109, September 2010.