Eliminating Reconciliation Cost in Secret Key Generation for Body-Worn Health Monitoring Devices

Syed Taha Ali, Vijay Sivaraman, and Diethelm Ostry

Abstract—Medical data collected by wearable wireless sensor devices must be adequately secured. A prerequisite for mass deployment of these secure systems is the capability of renewing cryptographic keys periodically without user involvement. Recent work has shown that two communicating devices can generate secret keys directly from measurements of their common wireless channel, which is symmetric but cannot be inferred in detail by an eavesdropper. These schemes may however yield mismatching keys at the two ends, requiring reconciliation mechanisms with high implementation and energy costs, unsuitable for resource-poor body-worn devices.

In this work we demonstrate a scheme for secret-key generation able to construct shared keys with near-perfect agreement, thereby avoiding reconciliation costs. Our specific contributions are: (1) we identify non-simultaneous probing of the channel by the link end-points as the dominant cause of channel measurement disagreement, (2) we develop a practical filtering scheme to reduce this disagreement, dramatically improving signal correlation between the two ends without affecting key entropy, and (3) we show that by restricting key generation to periods of significant channel fluctuation, we achieve near-perfect key agreement. We demonstrate in several representative body-worn settings that our scheme can generate secret bits with 99.8% agreement, and so yield near-perfect matching 128-bit keys approximately every half hour.

Index Terms—Body Area Networks, Secret Key Generation, Physical Layer Security.

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 personalised healthcare. A body area network typically consists of one or more small sensors mounted on the body to measure vital signs and communicate them wirelessly to each other and/or to a base-station (a fixed access point in the home or a portable device such as a mobile phone carried by the patient) for real-time analysis and response, and 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 mobile 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 is also critical, since these devices record and handle medical data which comes with stringent privacy and liability concerns. Some devices may also be actuators, for example delivering metered medications like insulin, and such devices present a security vulnerability with potentially serious medical consequences [5], [6].

The high energy and implementation costs of asymmetric cryptography precludes its use for encrypting medical data in a typical body-worn device, leaving symmetric (or shared key) encryption as the only viable option. The challenge lies in refreshing the secret keys shared by 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 [7] that users (such as the elderly) are often unaware of the need, or unable to configure keys of sufficient strength, or protect them 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 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 [8].

- *S. T. Ali and V. Sivaraman are with the School of Electrical Engineering and Telecommunications, the University of New South Wales, Sydney, Australia. E-mail: [taha, vijay]@unsw.edu.au*
- *D. Ostry is with the CSIRO ICT Centre, Marsfield, Sydney, Australia. E-mail: diet.ostry@csiro.au*

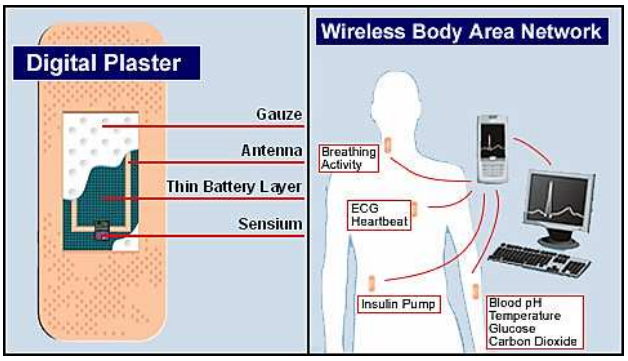

Fig. 1: The Sensium Digital Plaster and the associated body area network topology

Recent work such as [9], [10] has shown that it is possible to generate a shared secret over an unsecured wireless channel 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 in time and frequency to allow them 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) 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 [11] are proposed, where the two ends exchange messages to probabilistically identify mismatching bits.

In contrast to high bit rate key generation which is typically used for bootstrapping secure communications between two devices, we focus on low-data-rate patient monitoring applications that require only periodic key renewal. Pairwise temporal keys (or session keys) are recommended in the emerging IEEE 802.15 standard [12] for body area networks. In these applications a high bit generation rate is not essential; for example, if we assume that a 128-bit key needs to be renewed once every hour, as is recommended for WiFi [13], 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 and so eliminate reconciliation overheads which would consume precious computing and communication resources [14]. 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 so 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 transmissions. This significantly reduces radio usage, usually the most expensive operation in small sensor devices.

In this paper, we demonstrate shared secret key generation for the specific setting of a body-worn device communicating with a stationary off-body base-station. We propose a cost-effective scheme to eliminate key mismatch between the two endpoints. 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 the identification of the dominant cause of the observed channel mismatch during motion: the time delay between consecutive measurements by the two ends of the link. We present a theoretical bound on the mismatch, and validate it via experiments with body-worn devices comparing a representative office environment with 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 and thereby virtually eliminating key-bit mismatch. We show that an activity threshold can be adjusted to yield near-perfect key agreement by trading-off key mismatch against key generation rate.

For our threat model, we allow stationary passive eavesdroppers in the environment to sample the channel at the same time as the communicating parties with full knowledge of the key extraction algorithm and its parameters. 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, 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 uniquely constrained context of body-worn healthcare devices. We test our solution using off-the-shelf hardware, with one device worn on the human body in conditions approximating actual deployment of such devices and in typical usage environments such as an indoor office and a food court. Moreover, our scheme dispenses with

reconciliation and dedicated channel sampling, whilst generating high entropy secret bits at a usable rate of approximately 8 bits/min, with 99.8% bit agreement. At this rate, a usable 128-bit key is generated every 20 minutes. If a session key is renewed over a longer period, say 1 hour, the probability of generating a perfectly matching key at both endpoints using our mechanism can be up to 99.5%. Our scheme is lightweight, implementable on the current generation of body-wearable devices, and suitable for large-scale deployment in personalised healthcare systems.

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

2 Background

In this section we briefly describe secret key generation and prior research. We furthermore distinguish our approach from others in that we are able to dispense with key reconciliation in a bodyworn setting.

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 under ideal conditions their measurements 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. Alice and Bob can measure either of these representations 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 corresponds with the well-known Jakes uniform scattering model [15] 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 consumption of energy. An approach more suited to energy-constrained devices characterises the channel using the time-evolution of received signal strength, which fluctuates because of motion by the users or changes in the environment, as a source of shared information [9], [10].

In practice, asymmetric components appear in these channel measurements because of transceiver differences, random noise, the influence of motion, either of the parties or other elements of the environment, on the measurement procedure, 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 [9], [10], [16], although other suitable channel characteristics have also been studied [17], [18], [19].
2. *Quantisation:* The measurements are converted into a string of key bits. Approaches based on signal extrema [9], [10] and ranking [20] 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 [21].
4. *Privacy amplification:* The now matching keys are then strengthened by performing a transformation to increase key entropy and obscure any partial information an eavesdropper may have gathered during key reconciliation communications.

2.1.3 Reconciliation

We now consider the reconciliation phase to show that it incurs an unacceptably high cost in bodyworn devices, thereby motivating the study in this paper. Information reconciliation mechanisms have been developed mainly in the context of quantum cryptography [21], 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 minimise 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 would therefore be the probability of encountering an even number of errors, which can be expressed as:

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

Consequently, the probability P that a key of length K has no errors is

$$P = (1 - P_q)^{K/b} \quad (2)$$

For example, if there is as little as a 2% chance of a bit mismatching between the endpoints, for a block size $b = 8$ there is approximately a 15% chance of uncorrected errors in a key of length $K = 128$, and in this case the key will have to be generated again. 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.

The public exchange of parity (or any information about the bitstring) is also a security risk in that it leaks information that may make it easier to attack the key. The exact extent of an eavesdropper's advantage depends on the reconciliation protocol and the amount and nature of the information the two parties exchange. For instance, if Cascade is used to reconcile two bitstrings having a 1% bit mismatch, approximately 10% of the key bits are 'exposed' (as detailed in [22]) to the eavesdropper, and for a 10% bit mismatch, the number of exposed bits rises dramatically to $57 \sim 63\%$, which then requires a privacy amplification process to obfuscate the key.

Reconciliation protocols such as Cascade typically perform this parity check multiple times and shuffle the bit sequence in a coordinated way before each test. This incurs significant memory and transmission overheads (as documented in [14]), much more important for miniature sensor devices operating with constrained resources. Furthermore, reconciliation will also add to design complexity, of particular concern since these protocols will typically be implemented in hardware 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, quantiser 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

$$H(X) = -\sum_{i=1}^{n} p(x_i) \log_2 p(x_i) \quad (3)$$

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 [23] to estimate entropy. If the generated key successfully clears the battery of tests it ensures that the key cannot be distinguished from a random string using known statistical techniques, and is therefore suitable to be used in symmetric key-ciphers such as AES, and also lightweight block ciphers such as LED, XTEA, Piccolo, PRESENT, and CLEFIA that are more likely to be used on small resource-constrained devices.

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 as discussed next.

2.2 Prior Work

Prior work in secret key generation for **802.11 WiFi** considers both static and mobile cases. The authors of [9] 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. [10] 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 [21]) necessary, along with privacy amplification.

In [24], the authors consider key generation in **ultra wide-band 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 [25], the authors measure at a sequence of frequencies to estimate the spectrum and extract keys with agreement of over 97% in static deployments.

In [20], 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 quantisation. One of the endpoints must be moved continuously in a 'random' manner to induce signal fading fluctuations. This approach is extended in [16] 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 unique constraints and operating conditions. Channel variation is complex and unpredictable due to motion, shadowing effects of the human body and multipath propagation [26]. Body-worn devices were first considered in [27], where the authors simulated a near-body channel and derived an upper bound on secret bit rate of 4 bits/s due to inherent limitations on channel entropy. They do not describe an actual key generation process. We explored the tradeoff between secret key generation rate and bit agreement for body area networks in [28] and proposed a "zero reconciliation" scheme in [29].

2.3 Our Focus

Our emphasis is on minimising key generation costs due to the limited resources of body-worn devices. In contrast to earlier schemes, we eliminate 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 or environmental changes to create signal entropy which is harnessed for secret key generation. In this paper, we extend our previous work [29] aimed at eliminating reconciliation for body-worn devices. We report experimental results for additional user scenarios and provide basic guidelines on how to select the parameter values used in our scheme.

3 Understanding Disagreement

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

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 [30], illustrated in Fig. 2: (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*, 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 as signal paths change either due to movement on the part of the nodes or in the environment.

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 over the interval 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.

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

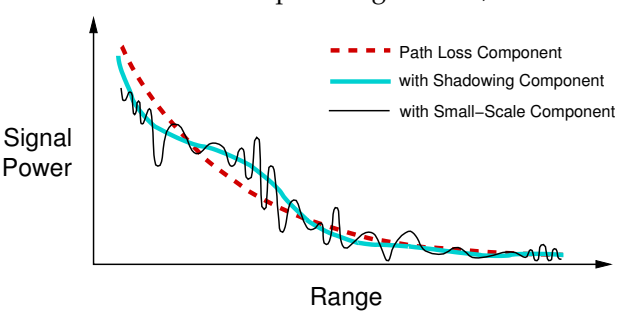

Fig. 2: Components of the received signal power

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 [31]. 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

$$f_{\max} = v \cdot (2/\lambda). \tag{4}$$

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

$$y(t) = (A/2) \sin 2\pi f_{\max} t \tag{5}$$

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

$$\begin{aligned} \Delta y &\approx dy/dt \cdot \Delta t \\ &\approx (A/2)\cos(2\pi f_{\max} t) 2\pi f_{\max} \Delta t \\ &\approx A\pi f_{\max} \Delta t, \quad \text{at } t = 0. \end{aligned} \tag{6}$$

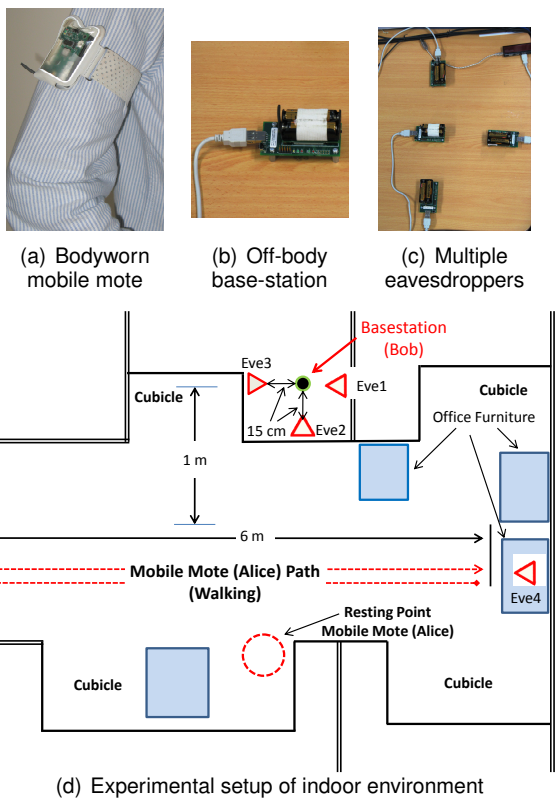

Fig. 3: Mobile node, base-station and experimental layout for indoor environment

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

$$\begin{aligned} \epsilon &= \pi f_{\max} \Delta t \\ &= 2\pi v \Delta t/\lambda. \end{aligned} \tag{7}$$

At an operating frequency of 2.4GHz for example (where $\lambda = 0.125$m) and a node velocity of $v = 1$m/s, a $\Delta t = 20$ms 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 change over the entire range from a minimum to a maximum during that interval. Since typical wireless sensor device radios today (e.g. the CC2420 [32]) take 10-20ms to probe the channel in one direction, this error can be significant in practice, causing mismatch between the two ends, as will be examined experimentally next.

3.2 Experiments in Indoor Environment and Anechoic Chamber

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

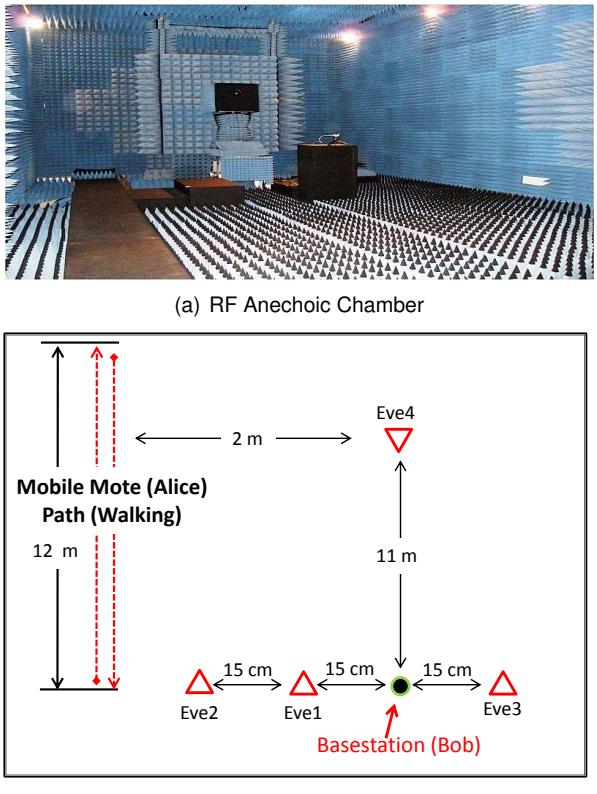

Fig. 4: Anechoic chamber and experimental layout

Our experiments used MicaZ motes running TinyOS and operating in the 2.4 GHz band. Their radios provide 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 averaging 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.

The layout for our *indoor environment* experiments, depicted in Fig. 3(d), show locations 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 in buildings around the anechoic chamber, but our results did not show evidence of interference. (It is relevant to mention here that efforts are underway to allocate spectrum specifically for bodyworn applications, to limit interference from other systems [33]).

The *RF anechoic chamber* is pictured in Fig. 4(a). All surfaces (floors, ceilings, walls) are covered in material that absorbs electromagnetic energy, thereby minimising 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.

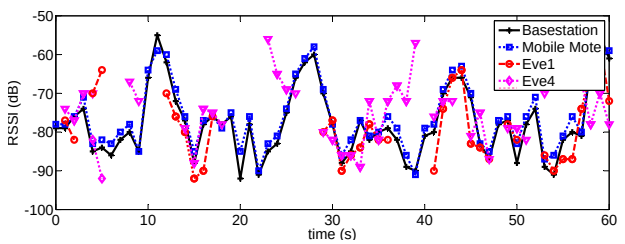

(a) RSSI trace, indoor office

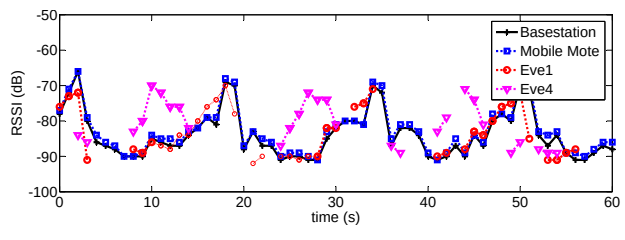

(b) RSSI trace, anechoic chamber

Fig. 5: Measurements comparing RSSI in the Indoor Office and in the Anechoic Chamber

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 private keys. However, we see 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 this discrepancy in RSSI more closely in Fig. 6 where a box plot depicts the variation in signal strength between the base-station and mobile node. 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, most of the discrepancy lies within an 8dB range (-6dB to $+2$dB), and when the signals are quantized, this results in significant bit mismatch i.e. reduced key agreement, at the two ends. The discrepancy is clearly much lower in the anechoic chamber, where it is almost completely concentrated at the median (1dB).

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

$$r = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \bar{X})^2}.\sqrt{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}} \quad (8)$$

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

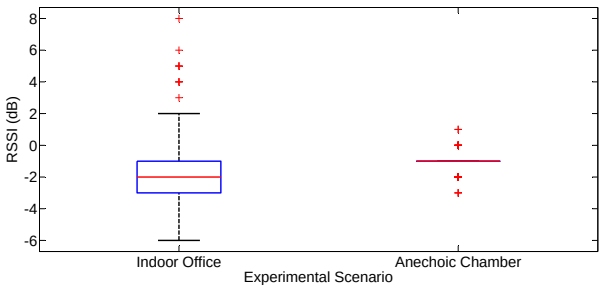

Fig. 6: Box plot highlighting the discrepancy in RSSI for both test environments

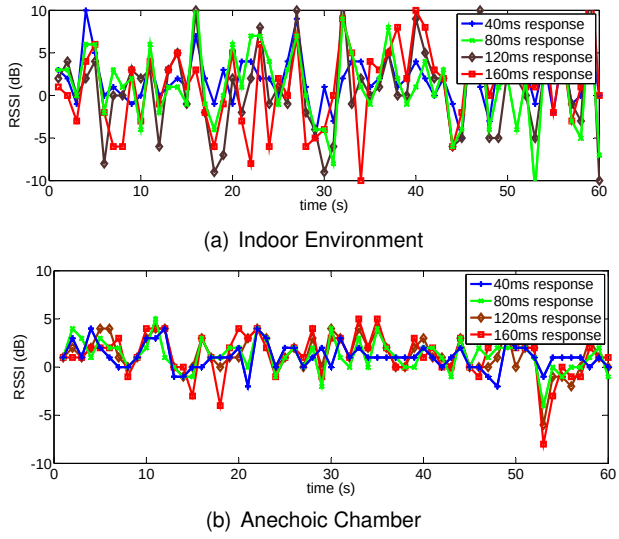

(a) Indoor Environment

(b) Anechoic Chamber

Fig. 7: Mismatch due to probing delay Δt

in RSSI measurements arising from differences in 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 discrepancies (which in turn leads lower key agreement) between the two communicating parties.

We can also validate experimentally that the discrepancy increases with increase in probing delay Δt. We configured the mobile node to acknowledge packet reception from a base-station several times successively 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 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 further (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 as probing delay increases, and that a 40ms probing delay in the indoor environment is equivalent to a 100ms probing delay in the multipath-free anechoic chamber in the sense of yielding a similar correlation of about 0.976.

These theoretical and experimental observations provide strong evidence that the discrepancy in channel measurement is predominantly due to the staggered sampling by the two ends of the link. In the

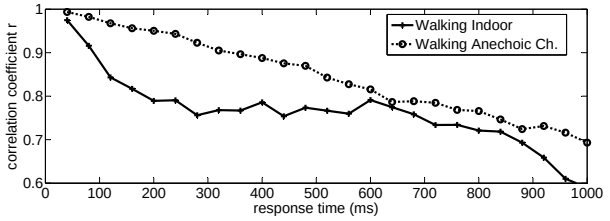

Fig. 8: Correlation coefficient r versus sampling delay

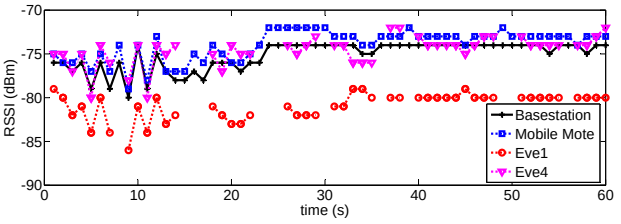

Fig. 9: Variation in RSSI for Resting Scenario

next section, we develop a novel means of reducing this discrepancy.

We wish to emphasise 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 stationary (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 the generation of mismatching keys, since the (uncorrelated) noise is amplified by the quantiser in the process of generating 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 minimise 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 [34] 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.

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 as $\epsilon \approx 100\%$. 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 [35] 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 [36], 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.

In the experimental work reported in this paper we select the parameters of the Savitzky-Golay filter to provide 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 [36, Eq. (11)]:

$$f_c \approx \frac{K+1}{1.6F - 3.6}, \tag{9}$$

where K is the polynomial order used by the filter, and F is the frame (window) size. We chose $K = 5$ (i.e. 5-th order polynomial) and $F = 9$ (for an impulse response half-length of 4), 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 emphasise 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 information about changes in the multipath with position that is 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 components arising from the effect of channel measurement delay on the small-scale fading, and which are primarily responsible for the disagreement between the two ends.

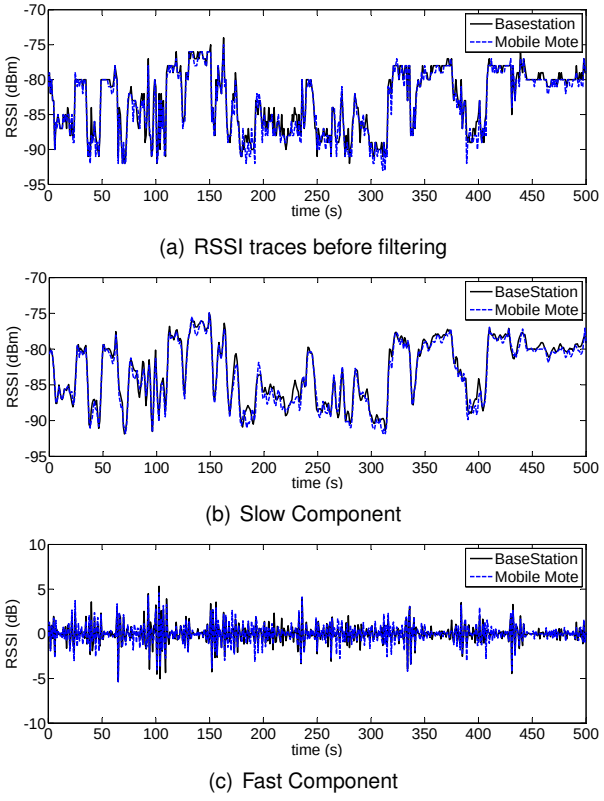

Fig. 10: Application of Savitzky-Golay filter

Comparing Fig. 10(b) with 10(a), we see that filtering visibly improves agreement between base-station and mobile node RSSI traces. 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 endpoints can be greatly improved by filtering the RSSI signals to attenuate 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 quantisation. We next propose a novel means of dealing with such effects by restricting key-bit generation to periods when channel fluctuation is not dominated by system noise effects.

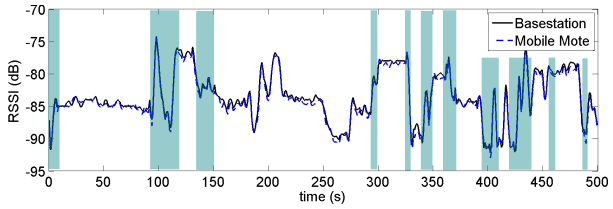

(a) Channel variation (slow component)

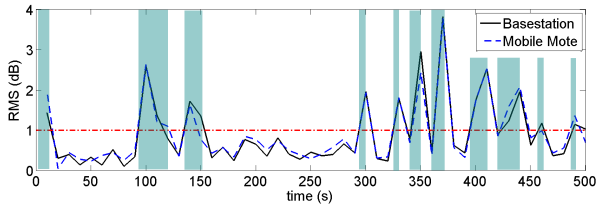

(b) Energy of fast component

Fig. 11: Region Selection for routine office activity

5.1 Dynamic Region Selection

When the channel exhibits significant fluctuations (e.g. when the subject is moving rather than resting or when the environment is dynamically changing), 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 [10], [16] 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 those 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 indoor office environment (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 root mean square (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 using the RSSI slow component.

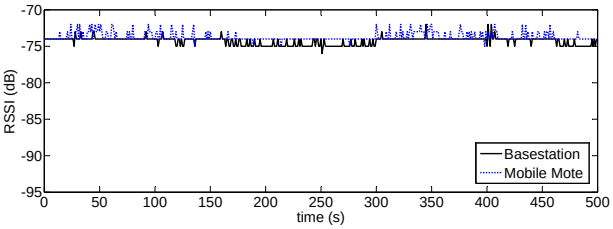

(a) Channel RSSI variation before filtering

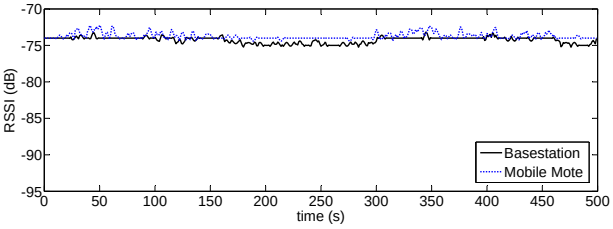

(b) Slow Component

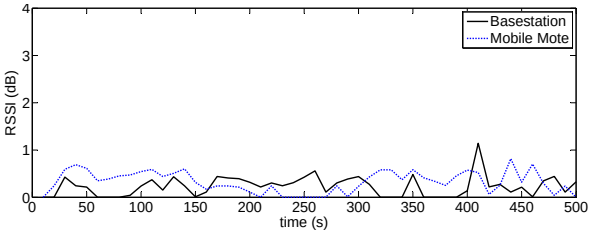

(c) Energy of Fast Component

Fig. 12: Energy of fast component when resting

5.2 Understanding Threshold θ

As we noted in Section 3.2, during resting and static scenarios where the half-duplex delay does not cause discrepancies, we observe a secondary source of key-bit discrepancy in small-scale fluctuations of RSSI caused by system noise. These fluctuations are uncorrelated between the two communicating parties, have small amplitude, and are a function of the particular radio hardware used at both ends. Quantisation of the RSSI measurements in the receivers can introduce further errors for small signal fluctuations. Discrepancies result if channel measurements made under these conditions are used to generate keys.

We can estimate a ceiling for the relative energy in these random components in a given system by measuring the channel variation in the static case. As shown in Fig. 12, when our system is in a completely stationary state, the channel RSSI varies by 1 to 2dB. The variation in the slow component is smaller due to the filtering used and has a variation of about 1dB from the average for both base-station and mobile node. The fast component energy (Fig. 12(c)) peaks at about 1.14dB for the base station and 0.82dB for the mobile mote. Based on these observations, it is reasonable to restrict key generation to periods when $\theta > 1.14$dB. In the next section, where we generate keys for θ varying from 0-3dB, we see this is a good approximation: we meet our intended target of 99.8% bit agreement when $\theta = 1.5$dB.

Parameter θ thus gives us a means to differentiate effectively between the static case, where channel variation is dominated by uncorrelated noise effects unsuitable for key generation, and regions of significant activity where variation is highly correlated at both ends and well suited for key generation.

5.3 Key Generation Flow

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 its parameters. However, we stipulate that Eve is separated from the two parties by a distance greater than one radio wavelength (~ 12.5 cm for the 2.4GHz band), and thereby forced to measuring a different multipath channel. We do not consider here the issue of initial trust between base-station and mobile node, nor active attackers engaged in jamming and packet injection.

The key generation mechanism runs as a background process to normal device operation, with the process flow depicted in Fig. 13 which identifies 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 (polynomial order K and frame size F are preconfigured) which outputs the "slow component". The "fast component" is obtained by subtracting the slow component from the original signal, 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 quantiser for bit generation. We note here that all of these operations are linear and can easily be implemented in hardware.

Our research does not develop a new quantiser. Instead, we use a basic single-bit quantiser, taken from [9] and refined in [10], 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. 14. For each block, two threshold values are calculated:

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

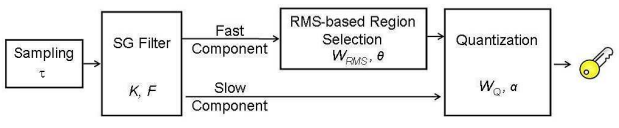

Fig. 13: Flow chart of key generation process

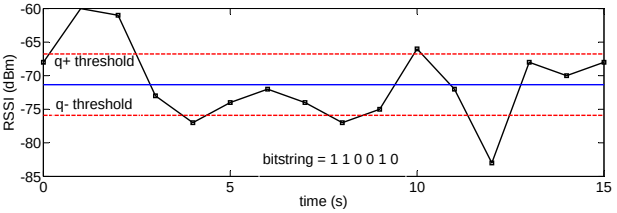

Fig. 14: Quantisation Process

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 quantiser performance to balance between bit generation rate and mismatch. For our purposes, we fixed the window size to $W_Q = 5$ and set $\alpha = 1$, consistent with prior work [10].

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

6 Results and Analysis

We tested our key generation mechanism in the office space in Fig. 3(d) and in a busy public space, i.e. a food-court on the University of New South Wales campus. 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 fetching items from other cubicles. In the third experiment, *Dynamic Environment*, performed in the food court, the intention was to keep the two communicating parties stationary relative to each other, and use customer traffic in the surrounding environment to cause the channel fluctuations needed to generate secret-key bits. The subject was seated at a table about 6m away from the base-station, almost double the separation in the office environment. This experiment was conducted during lunch hour when there was a maximum of pedestrian traffic. Trace data was collected from each experiment for about 40 minutes and our key generation scheme was applied offline to assess its performance with different parameter settings.

The base-station and mobile node were again micaZ motes and sampled the channel once per second ($\tau = 1$)/s. We use the Matlab implementation of the

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

Signal quantised	Key Agreement (%)	bit rate (bit/s)	Eve1 Key Agreement (%)	Eve2 Key Agreement (%)	Eve3 Key Agreement (%)	Entropy
unfiltered	98.40	0.205	47.11	46.48	47.34	0.9970
filtered, $\theta = 0$	97.91	0.244	50.89	51.00	51.03	0.999
filtered, $\theta = 0.5$	99.08	0.222	50.49	50.64	50.79	0.999
filtered, $\theta = 1$	99.74	0.181	50.66	50.82	50.97	0.998
filtered, $\theta = 1.5$	99.83	0.141	50.22	50.56	50.51	0.999
filtered, $\theta = 2$	99.88	0.101	50.14	50.16	50.10	0.998
filtered, $\theta = 2.5$	99.92	0.065	50.35	50.22	50.27	0.998
filtered, $\theta = 3$	100	0.037	49.81	49.15	49.95	0.998

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

Signal quantised	Key Agreement (%)	bit rate (bit/s)	Eve1 Key Agreement (%)	Eve2 Key Agreement (%)	Eve3 Key Agreement (%)	Entropy
unfiltered	95.53	0.139	46.26	46.80	47.60	0.9971
filtered, $\theta = 0$	93.06	0.197	48.40	47.88	48.30	0.999
filtered, $\theta = 0.5$	98.41	0.132	48.39	47.79	48.49	0.999
filtered, $\theta = 1$	99.41	0.086	48.26	47.74	48.41	0.999
filtered, $\theta = 1.5$	99.80	0.057	47.81	46.92	48.01	0.999
filtered, $\theta = 2$	100	0.036	47.38	45.79	47.69	0.999
filtered, $\theta = 2.5$	100	0.024	47.35	44.73	47.58	0.999
filtered, $\theta = 3$	100	0.015	46.54	43.85	47.05	0.999

TABLE 3: Effect of varying threshold θ on key generation performance metrics for *Dynamic Environment* scenario

Signal quantised	Key Agreement (%)	bit rate (bit/s)	Eve1 Agreement (%)	Eve2 Agreement (%)	Eve3 Agreement (%)	Entropy
unfiltered	99.86	0.23	53.22	52.48	52.84	0.9968
filtered, $\theta = 0$	99.92	0.250	50.44	50.99	49.80	0.999
filtered, $\theta = 0.5$	99.92	0.246	50.48	50.89	49.68	0.999
filtered, $\theta = 1$	99.89	0.180	48.57	49.89	48.68	1.000
filtered, $\theta = 1.5$	100	0.068	45.75	47.21	45.16	0.995
filtered, $\theta = 2$	100	0.015	44.16	42.86	37.66	0.985
filtered, $\theta = 2.5$	100	0.002	50.00	37.50	37.5	0.954
filtered, $\theta = 3$	100	0.001	33.33	33.33	33.33	0.918

Savitzky Golay filter where polynomial order, $K = 5$, and frame size $F = 9$. The energy of the fast component is computed over a window size $W_{RMS} = 10$. The quantiser window size, $W_Q = 5$, where W_Q is chosen to be a factor of W_{RMS} (to ensure that the bitstrings synchronise, in case either party encodes extra bits), and quantiser $\alpha = 1$. The value of the RMS threshold, θ, is then varied in increments of 0.5 from 0 to 3dB to note its effect on performance.

High Activity: Table 1 shows the percentage of key bits that agree for different energy threshold settings. It is evident that filtering by itself does not significantly improve signal correlation between the two ends (the agreement actually decreases a very small amount from 98.40% to 97.91%). This can be explained by the fact that the subject is not in a constant state of motion, unlike in our walking experiments earlier, and that during quiescent periods, the quantiser amplifies uncorrelated random noise (as explained in Section 4) which is likely to cause key disagreement. However there is a marked improvement in key generation rate when region selection is applied to restrict key generation to regions with at least θdB energy in the fast component: a threshold setting of $\theta = 0.5$dB improves key generation to over 99%, and at $\theta = 1.5$dB, over 99.8% of the bits match.

This improved agreement comes however at the cost of reduced bit rate, which decreases from 0.205 to 0.141 bits/s. Region selection reduces bit generation rate, because with increasing threshold, a progressively smaller proportion of the signal is available for quantisation. This trade-off is illustrated in Fig. 15(a), 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. At this generation rate, a usable 128-bit key is generated approximately every 15 minutes.

Low Activity: Table 2 shows that agreement of keys generated from the raw unprocessed signal is quite low at around 95%. This can be attributed to longer quiescent or low-motion periods during this experiment where the subject just sits at his desk. Again, filtering combined with region selection has 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 at 99.8%, which meets our intended bit agreement target. Key generation rate similarly decreases from 0.139 to 0.057 bits/s, and a usable 128-bit key can be constructed in about 35 minutes. The trade-off is shown in Fig. 15(b).

Dynamic Environment: In this scenario the base-station and mobile node are stationary, and RSSI variation is caused primarily by changes in the environment, in particular by the motion of people walking

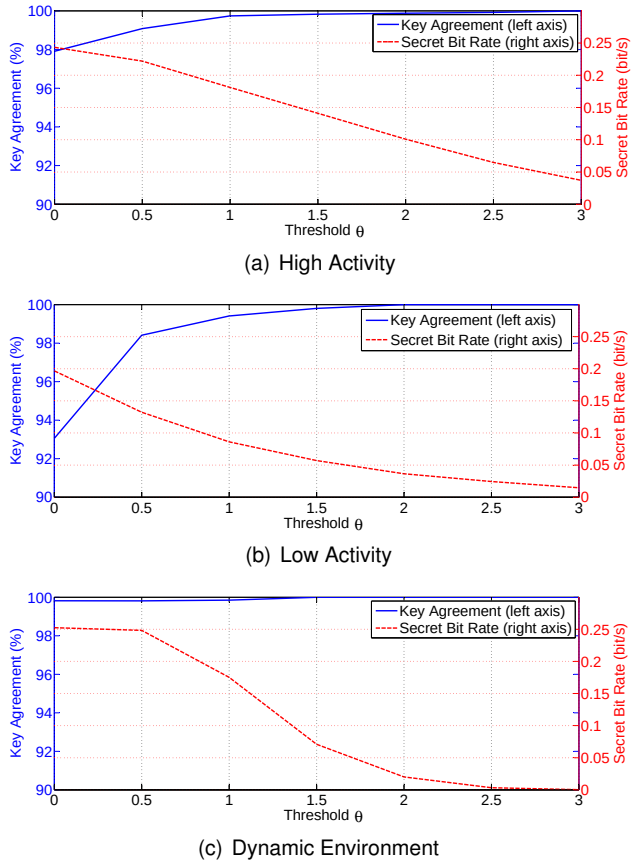

Fig. 15: Key agreement vs. secret bit rate for varying region selection threshold θ

along paths mainly located between the mobile mote and base-station. These people were moving at typical speeds, but given the large amount of traffic dispersed throughout the room, multiple signal paths were affected at the same time, with a correspondingly greater magnitude of change per unit time compared to the other scenarios (where a single person, the mobile subject, is responsible for the channel fluctuation). This had the notable advantage that the unprocessed signal was already adequate fluctuating for key generation, yielding an agreement of 99.86% (which meets our desired target) at a rate of 0.234 bits/s. Applying region selection with $\theta = 1.5$dB improves agreement to 100% at a rate of 0.068 bit/s, giving a usable 128-bit key in approximately 30 minutes. The trend in variation of key agreement and bit rate with changing threshold is depicted in Fig. 15(c).

These preliminary results allow us to select a threshold θ value which can give good performance across a wide range of different environments and activities. For our radio hardware and the scenarios we investigated (which might be considered representative of typical use), 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 to 0.141 bits/s, i.e. it would take 15 to 35 minutes to generate a usable 128-bit key, which we believe is fast enough for purposes of key renewal. If typical session key lifetime is approximately 1 hour, the probability of having a valid new key perfectly matching at both endpoints therefore varies from 93.5 to 99.5% depending on the user's activity.

The high key entropy seen in all cases (> 0.99) (Column 6), and the keys' passing the NIST *approximate entropy* test [23] confirm that the filter retains a sufficient component of the essential randomness arising from motion in a multipath environment.

Tables 1, 2, and 3 also show the percentage of matching bits that each of the eavesdroppers generate by passively listening to the channel. Eavesdropper agreement hovers near the ideal 50% for almost all cases, indicating that their chance of guessing if a generated bit is correct or not is no better than an unbiased coin toss. For the last case, *Dynamic Environment* and $\theta = 3$dB, the sharp drop in eavesdropper agreement (to 33%) is due to the fact that too few key bits were generated by all parties in that run (only 3) to have statistical significance.

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 primary cause of key mismatch: the delays in measuring the channel in both directions due to half-duplex radio operation. 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. Furthermore, we noted that these discrepancies are concentrated in the rapidly-varying component of the channel RSSI trace. Second, we showed that this rapidly-varying 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 key reconciliation, and incrementally generates high entropy key bits at a rate suitable for key renewal. Experimental results show that it takes 15 to 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 to 99.5% chance of generating a perfectly matching secret-key which can be used by common ciphers such as AES, LED, PRESENT, etc. which is a promising result.

In future work, we intend to extend this approach to wholly on-body communications, and secure the wireless links between bodyworn sensors, portable base-stations, and/or medical implant devices.

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] C. Li, A. Raghunathan, and N.K. Jha. Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system. In *IEEE Healthcom 2011*.

[6] W. Burleson, S.S. Clark, B. Ransford, and K. Fu. Design challenges for secure implantable medical devices. In *ACM/EDA/IEEE DAC 2012*.

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

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

[9] 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.

[10] 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.

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

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

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

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

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

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

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

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

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

[20] 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.

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

[22] T. Calver. An Empirical Analysis of the Cascade Secret Key Reconciliation Protocol for Quantum Cryptography. Master's thesis, Air Force Institute of Technology, September 2011.

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

[24] 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.

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

[26] D. Smith, L. Hanlen, A. Zhang, D. Miniutti, D. Rodda, and B. 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.

[27] 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.

[28] S. T. Ali, V. Sivaraman, and D. Ostry. Secret Key Generation Rate vs. Reconciliation Cost using Wireless Channel Characteristics in Body Area Networks. In *IEEE TrustCom-10*.

[29] S. T. Ali, V. Sivaraman, and D. Ostry. Zero Reconciliation Secret Key Generation for Body-worn Health Monitoring Devices. In *ACM WISEC'12*.

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

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

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

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

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

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

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

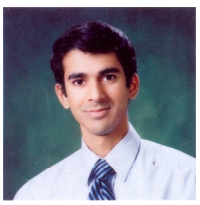

Syed Taha Ali did his BSc. (Eng) from GIK Institute of Engineering Sciences and Technology, Pakistan, in 2002 and his MS and PhD in Electrical Engineering from the University of New South Wales, Australia, in 2006 and 2012. His research interests include wireless sensor networks, network mobility, and security.

Vijay Sivaraman received his B. Tech. from the Indian Institute of Technology in Delhi, India, in 1994, his M.S. from North Carolina State University in 1996, and his Ph.D. from the University of California, Los Angeles in 2000. He has worked at Bell-Labs and a Silicon Valley start-up manufacturing optical switch-routers. He is now Associate Professor at the University of New South Wales, Australia. His research interests include optical networking, packet switching and routing, Quality of Service, and wireless sensor networks.

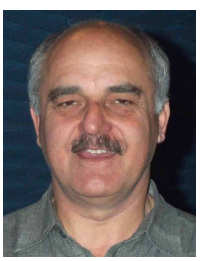

Diethelm Ostry is a Research Scientist in the Network Technologies Laboratory, Information and Communication Technology Centre, CSIRO Australia. His recent research interests have been in the areas of network traffic characterisation, security in wireless sensor networks, and cooperative communications.