Securing First-Hop Data Provenance for Bodyworn Devices using Wireless Link Fingerprints

Syed Taha Ali, *Member* Vijay Sivaraman, *Member* Diethelm Ostry, *Member*
Gene Tsudik, *Senior Member* and Sanjay Jha, *Senior Member*

Abstract—Wireless bodyworn sensing devices are fast becoming popular for fitness, sports training and personalized healthcare applications. Securing data generated by these devices is essential if they are to be integrated into the current health infrastructure and employed in medical applications. In this paper, we propose a mechanism to secure data provenance for these devices by exploiting spatio-temporal characteristics of the wireless channel that these devices use for communication. Our solution enables two parties to generate closely matching 'link fingerprints' which uniquely associate a data session with a wireless link such that a third party can later verify the details of the transaction, particularly the wireless link on which the data was transmitted. These fingerprints are very hard for an eavesdropper to forge, they are lightweight compared to traditional provenance mechanisms, and enable interesting security properties such as accountability, non-repudiation, and resist man-in-the-middle attacks. We validate our technique with experiments using bodyworn sensors in scenarios approximating actual device deployment, and present some extensions which reduce energy consumption. We believe this is a promising first step towards using wireless-link characteristics for data provenance in body area networks.

Index Terms—Body Area Networks, Data Provenance, Physical Layer Security.

I. INTRODUCTION

Body area networks (BANs) is an emerging technology paradigm anticipated to revolutionize the healthcare domain and significantly reduce soaring national health expenditures. Small unobtrusive sensors worn on the body allow for mobility, remote monitoring, and reduce the burden on hospital and professional staff. This technology has already found popular application in sports, fitness training, and lifestyle monitoring. Examples include Nike+ FuelBand [1], Fitbit Flex [2], Toumaz Sensium Digital Plaster [3], and the Natalia Project [4]. Several companies, including Apple [5], are reportedly innovating in this field, and ABI Research [6] predicts that shipments of disposable wireless sensors are expected to reach 5 million by 2018.

A typical body area networks topology consists of bodyworn sensors which communicate over a wireless link with a handheld device or an off-body basestation which forwards the data to an online database to be accessed and analyzed by

Syed Taha Ali, Vijay Sivaraman and Sanjay Jha are with the University of New South Wales, Kensington, Sydney, NSW, 2052 Australia, e-mail: taha, vijay, sanjay@unsw.edu.au.

Diethelm Ostry is with CSIRO ICT Centre, Vimera Rd., Marsfield Australia, NSW, 2122, Australia, email: diet.ostry@csiro.au.

Gene Tsudik is with University of California, Irvine, CA 92697, United States, email: gtsudik@uci.edu.

professionals. A sensor may communicate directly with the basestation (as in a star topology) or it may forward the data to it over multiple hops using other sensors. Since sensors are severely resource-constrained, they cannot use traditional cryptographic solutions due to the high overhead, and are the weakest link in this architecture. However, security is clearly needed because these devices deal in personal medical data, wrongful disclosure and tampering of which can result in serious ethical and legal implications. Developing lightweight security solutions for these devices is therefore a popular research area.

Numerous methods have been proposed for confidentiality, integrity, and authenticity of sensor data. However, for these devices to integrate successfully into the healthcare infrastructure and for patients and medical professionals to trust this data, further guarantees must be provided, such as contextual information about the data, which may include information about sensor-patient association, data-device association, and which parties handled the data. This metadata falls in the purview of *data provenance*. To this end our provenance approach specifically generates a secure verifiable proof of data transactions occurring between sensor device and basestation.

Consider the case of a user, Alice, who has had a heart attack and is informed by her insurance provider that her insurance rates will be reduced if she gives up smoking. To ensure she complies, Alice is given a wearable sensor device to monitor her for a trial period. The sensor periodically sends readings to her smartphone which forwards them to an online database. Thwarting this mechanism to secure benefits is easy: Alice could easily hack into the phone and forge her readings. Or she could replay previously recorded readings to cover up smoking episodes. Or, in the identity transference attack, she could even affix the sensor to a non-smoker friend for the duration of the trial without anyone finding out. Traditional mechanisms to protect against such attacks rely on extensive usage of public-key cryptography which is compute and energy intensive.

This paper proposes a data provenance technique for bodyworn sensors. Provenance facilitates data trustworthiness of the data, which is a critical factor especially in data forensics. In Alice's case, it would be useful to reliably determine certain information about the data such as, the common sensor data off-load points. It would enhance trustworthiness of the medical readings to confirm that Alice's sensor's first point of contact is usually her personal mobile phone, home WiFi access point, a basestation in her office, or her gym. Forensics investigators should be able to check data associations, verify

context, identify faults and, in the event of an incident, assign liability.

Prior work in provenance for sensor devices (reviewed in the next section) typically relies solely upon cryptographic mechanisms which may prove very energy intensive [7]. In contrast, we suggest an information-theoretic approach: we propose that wireless channel characteristics between the sensor and the basestation be used to generate 'link fingerprints'. Characteristics of the wireless link, such as radio signal strength or signal phase shift, are unique to the two communicating parties, very difficult for an eavesdropper to forge, and can be leveraged to provide a shared and provable record of data sessions between two devices. Public-key operations, such as digital signatures, are used sparingly only to authenticate link fingerprints on a per session basis.

When a sensor and a basestation communicate, their routine data messages can be used to sample the wireless link, allowing each party to generate a unique very closely matching signature or 'link fingerprint' (usually RSSI is used for this purpose). This is similar, conceptually, to a Diffie Hellman key exchange, except that it is done at a fraction of the processing cost, and computed over the course of the entire data exchange.

If both parties digitally sign the data they exchange and their corresponding link fingerprints, this would authenticate the session, and allow for later (off-line) verification of the sensor-basestation association for the transaction, effectively confirming that said data was transmitted over that particular link. This process is **secure** since the fingerprints cannot be forged. It is also **lightweight** compared to alternative solutions that rely solely on cryptography. Unlike existing provenance solutions, our scheme also provides **accountability**, i.e. the wearer of the sensor can verify that provenance information has not been tampered with or hacked. This is especially applicable in scenarios involving remote programming, where data flow is from basestation to sensor device. For example, considering bodyworn actuator devices (such as insulin pumps), Alice (or her doctor) may choose to reprogram her bodyworn device, and, the signed link fingerprint would serve as **non-repudiable** evidence of that operation.

Our contributions are:

1) a data provenance protocol using wireless channel characteristics to generate link fingerprints.
2) experimental results confirming that this protocol can generate unique and near-perfect matching link fingerprints for typical data exchanges between bodyworn sensor and off-body basestation.
3) optimization mechanisms that significantly reduce memory and transmission overheads in handling link fingerprints, making the proposed protocol feasible for resource-constrained devices.

Our results indicate that, in a typical usage environment, two parties can generate a 128-bit link fingerprint approximately every $10 - 15$ minutes. We believe this is a promising first step in using wireless link characteristics to enable secure data provenance.

This paper is organized as follows: Section II covers prior work in this domain. In Section III, we summarize research in exploiting radio channel state variation to build security primitives. The link fingerprint protocol is described in Section IV. The fingerprint generation technique is experimentally validated in Section V, optimization mechanisms are presented in Section VI. The paper concludes in Section VII.

II. PRIOR WORK

We consider **provenance** to be a record of the origin and evolution of data within a system. On a computer, provenance may consist simply of a record of processes involved in system events pertaining to the data, such as creation, access, modification, etc. In a digital domain provenance may include a record of paths the data took, and a log of any remote actions performed on it. This is vital for digital forensics: with the surge in computer crime, provenance is critical in reconstructing incidents and assigning liability. This also motivates the need for *securing* provenance (distinct from generating it), as discussed in [8].

The **granularity** of provenance varies with application requirements and device capability: [9] makes the case for 'high-fidelity' provenance, compiled at the kernel level, enabling very detailed forensics analysis. On enterprise networks, administrators can log file and system operations in detail. On resource-constrained devices, however, a digital signature or a timestamp association may have to suffice. Additionally, provenance need not be binary: especially in distributed environments, such as large multi-hop sensor networks, it may be more practical to express confidence in sensor data using a probability value [10] or a trust score [11] [12].

In **body area networks**, provenance has mostly been limited to verifying data-sensor and data-patient associations. Chowdhury et al. [13] survey existing research on associating sensor data with the human subject, and consider several authentication techniques, typically relying on frequent use of cryptographic protocols, trusted third parties, and additional sensor node capability (e.g. biometrics readers). [14] amortizes digital signatures for bodyworn devices, enabling a secure and irrevocable binding between patient data and the originating device. [15] proposes binding patient data to the subject's own unique vital signs readings (real-time ECG and accelerometer data), enabling user authentication in a continuous manner. However, compared to our solution, none of these schemes address the data path or link association between two parties, and some of them require extra hardware capabilities.

Radio fingerprinting techniques [16] can identify a transmitting party (with up to 70% probability) by examining its radio signal. However, they require a strictly stationary deployment, specialised sampling hardware and a database to train the system. This approach is not scalable to multi-hop networks, and a sophisticated attacker may even forge legitimate radio fingerprints.

Identifying the links that data traverses in a wireless network is examined in [17], which proposes a provenance mechanism where intermediate nodes in a **multi-hop sensor network** use Bloom filters to imprint path information on transit packets such that the basestation can verify the path of each packet. This has the advantage that it can identify malicious nodes. However, it relies on a trusted infrastructure and does not protect against collusion or man-in-the-middle attacks.

A similar concept to our work is that of **secure routing protocols**: for instance, in the Secure Ad-hoc On-demand Distance Vector (SAODV) [18], routers in a mobile ad hoc network digitally sign the headers of all routing information packets so that devices in the network are able to verify the route the packet has taken. However, this results in intensive overheads due to the signing operations and transmission costs.

The distinguishing feature of our approach is the generation of provenance between two parties, on a *per session* basis, thereby minimizing the use of expensive cryptographic operations and communication overhead, and still providing strong security guarantees. This makes our approach ideal for resource-constrained bodyworn devices. In most prior work, digital signatures and/or encryption is used on a per packet basis, whereas in our case, these operations are performed only after a usable link fingerprint is derived, approximately every ten to fifteen minutes (after every 900 data packets). The link fingerprint is generated using only linear operations by exploiting wireless channel characteristics during the data transmission. We present next an overview of this channel-based technique and describe briefly how researchers are deploying it for a variety of applications, including secret-key agreement, authentication, and intrusion detection.

III. SECURITY FROM WIRELESS CHANNEL CHARACTERISTICS

There has been considerable interest recently in using the wireless physical channel between two devices to construct security primitives. The theory underpinning this approach is as follows: the wireless channel between two communicating parties, Alice and Bob, is intrinsically symmetric, i.e. if Alice and Bob were to use identical transceivers and antennas, and transmit identical signals, they would also receive identical signals. As described by the reciprocity property of electromagnetic communication, radio signals propagate over an identical set of multiple paths in both directions, thereby experiencing identical gains, phase shifts, and delays. Alice and Bob can individually measure parameters (such as radio signal strength, phase shift, angle of arrival, etc.) describing the effects of these paths on the signal, and in ideal conditions, barring interference and noise, these measurements will agree.

The wireless channel is also highly sensitive to location and spatio-temporal changes. Jakes fading model [19] states that the wireless channel decorrelates rapidly over distances of approximately half a wavelength, and, for a distance greater than one wavelength, the channels may be assumed to be independent. The implication is that motion on the part of Alice or Bob or in the environment will cause significant variation in the time-evolution of channel characteristics. Furthermore, if an eavesdropper, Eve, is located a distance greater than one radio wavelength, she will effectively be measuring a different spectrum, and will be unable to access Alice and Bob's shared measurements. Alice and Bob may therefore use their channel measurements as a source of shared entropy for security benefits.

We support this theory with experimental evidence in Section V. The effectiveness of this technique, however, has been demonstrated in the literature across a wide range of platforms and a variety of applications. In one of the earliest practical demonstrations in this area, the authors in [20] use software defined radios operating in the 400-500 MHz band and suggest methods for two-party authentication, secret-key agreement and the secret-key dissemination. Secret-key agreement is the most popular application and schemes have been devised for UWB communications [21], Bluetooth [22], and WiFi networks [23] [24] [25]. In[24], the authors specifically undertake a detailed investigation of shared channel state for various environments (including an underground tunnel, a cafeteria, and an outdoor lawn), different activities (walking and riding a bike), and also a scenario where Alice and Bob communicate using different radio hardware.

Considering bodyworn sensing devices specifically, extensive experimental campaigns have been conducted in recent years. In [26], the authors model channel fading for bodyworn devices in the 2.45 GHz band, whereas [27] focus on the 2.4 GHz and 900 MHz ISM and the 402 MHz medical implant communications band. The channel variation is observed to be complex and unpredictable due to small-scale fading caused by motion, multipath propagation, and the shadowing effects of the human body. In [28], the authors examine the feasibility of secret-key agreement using the radio channel as modelled by the IEEE 802.15.6 Task Group on body area networks.

Some research efforts (including ours) approximate real deployments using 802.15.4 low-end sensing devices mounted on the body, e.g. TelosB [29] [30] and MicaZ motes [31] [32] [33]. In [29], the authors propose a solution allowing multiple sensors worn by a subject to authenticate each other, and in [30], to undertake authenticated secret-key agreement. The authors in [31] decompose the wireless channel into fast and slow components, enabling the user to configure secret-key generation as per his requirements. [32] suggest the use of filtering techniques to reduce the effects of noise and asymmetric components in the channel profile and further improve the correlation between Alice and Bob for improved secret-key agreement. [33] perform extensive experiments (in an office space and an anechoic chamber) to quantify the effects of small-scale fading on secret-key agreement and propose a solution which restricts the key generation process to periods of high channel variation.

Other security applications that leverage the entropy of the shared channel state include intrusion detection [34], location distinction [35], secure pairing [36], proximity-based authentication [37], and identifying spoofing and Sybil attacks [38] [39]. However, to the best of our knowledge we are the first to exploit shared wireless channel characteristics for data provenance. We describe next a low-cost protocol, we highlight its novel security properties, and demonstrate with experiments its applicability for bodyworn monitoring devices.

IV. PROTOCOL FOR LINK FINGERPRINTS

Recall the earlier example of Alice that was used to motivate the need for associating a sensor with the basestation for the duration of a data transaction. We use the 'channel signature' to uniquely fingerprint the Alice-Bob link and associate it with
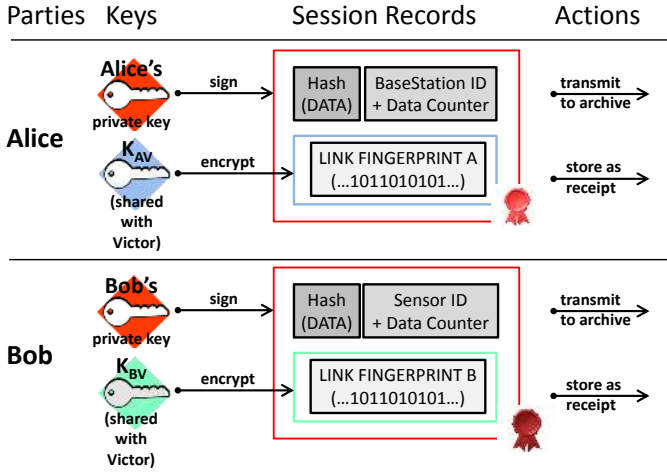
Fig. 1: Protocol for Alice and Bob

Fig. 2: Protocol for Victor

the data they exchange, such that a verifying party, Victor can later verify this association. The fingerprint is a unique bitstring of pre-configured length, generated individually by Alice and Bob sampling their common wireless link. Details of how both parties generate the fingerprints are provided below.

A. Sensor Device and Basestation

The protocol for the bodyworn sensor device and base-station, denoted as Alice and Bob respectively, is depicted in Fig. 1. It is executed at the conclusion of the transaction, after all data messages have been exchanged between the pair, and they have generated their respective link fingerprints. The Hash(DATA) field in this case refers to a hash operation performed over all the data the two parties have exchanged.

It is essential that Alice and Bob encrypt their fingerprints to keep them private from unauthorized parties. The fingerprint is a unique characterization of the spatio-temporal channel variation which is correlated at both ends of the wireless link. If this were transmitted in plain sight, an attacker could copy it and claim the link association with itself, raising confusion with regard to the actual data offload point. Alice and Bob should also not be able to view each other's fingerprint: if either of them were maliciously inclined, they might share the fingerprint with an attacker or use the fingerprint and its signature binding to try and mount a replay attack. An easy way to protect the fingerprint is for each party to encrypt its own with a key shared only with Victor (in this case, K_{AV} for Alice and K_{BV} for Bob).

Once the fingerprint is encrypted, it is bundled with a hash digest of the data and session identifiers (timestamp or counter value, identity of the devices, etc.), into a *session record* which is digitally signed, and transmitted to the database. The signature ensures both parties commit to the data and the link association.

Furthermore, each party can retain a copy as a *receipt* for the transaction, enabling system-wide **accountability**. In existing provenance solutions, sensors usually offload trust on to the network or third parties ([11] and [10]), and this allows for scenarios where trusted insiders may tamper with the provenance record without detection.
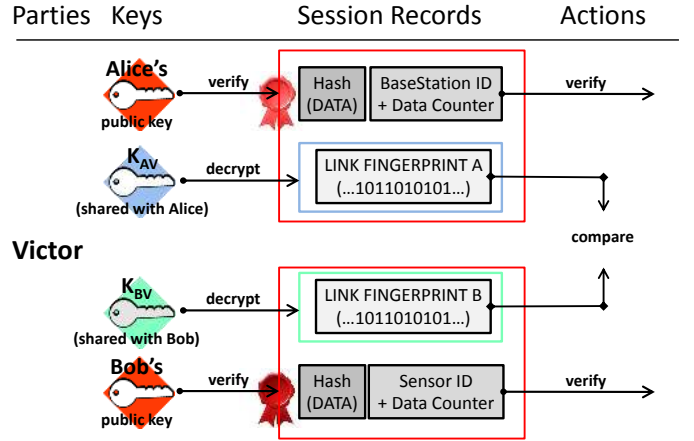
B. Verifier

Fig. 2 depicts the verification process for Victor: as part of a subsequent review or audit, Victor could revisit archived session records, verify the digital signatures and session identifiers for the data items of interest, decrypt Alice and Bob's link fingerprints using the individual symmetric keys and check that they match. If there is a very high similarity between fingerprints, Alice and Bob used the wireless link to communicate that particular data item.

Furthermore, if Alice were to act maliciously and alter the record of her transaction with Bob (by tampering with the fields in their session record), she would not be able to tamper with Bob's record of the session (due to his digital signature), and Victor would easily identify the mismatch.

C. Discussion

One important property of the proposed scheme is that the link fingerprint mitigates man-in-the-middle attacks. If Eve interposes herself between Alice and Bob, the legitimate Alice-Bob communication would span two different wireless links (Alice-Eve and Eve-Bob) and Alice and Bob finger-prints would markedly differ with very high probability. This property can be used to augment other locationing protocols vulnerable to man-in-the-middle attacks. We briefly consider two examples:

Man-in-the-middle attacks in wireless networks have been usually addressed using **distance-bounding protocols**. A distance bounding protocol [40] enables a party, Alice, to determine an upper bound on the physical distance to a communicating party, Bob. A common usage scenario is that of a person with a badge (access card) running a wireless identification protocol to gain access at the entrance to a building. The access control system must verify that the badge running the identification protocol is in the immediate vicinity.

In this protocol, Alice typically issues a challenge to Bob and measures timing delay between sending the challenge and receiving the response. The upper bound on distance to Bob is the round-trip time divided by twice the speed of light, since RF waves travel at the speed of light. State of the art distance bounding protocols (such as [41]) focus on minimizing

processing delay at Bob's end, enabling it to respond faster to Alice's challenge, thereby reducing uncertainty in distance estimate. Whereas distance bounding does not prevent the man-in-the-middle attack, it makes it much harder by forcing the attacker to be in the immediate physical proximity of Alice and Bob. If a secure distance bounding protocol were to be used in conjunction with link fingerprints, Alice would have two strong guarantees, (1) that she is communicating directly with Bob with no intermediary in between, and (2) Bob is in her immediate vicinity.

A similar higher level primitive is **location proofs**, introduced by Saroiu et al. [42]. In a typical usage scenario, a mobile device, seeking to prove its location may request a component of the local wireless infrastructure (such as a WiFi access point or a cell phone tower) to issue it a location proof. Location proofs enable several potential applications, such as location-restricted content delivery, store discounts and loyalty schemes, and fraud reduction in online auctions. The proof itself consists of timestamped metadata digitally signed by the wireless access point. Certain applications, such as voting registration or generating alibis for police investigations, may require more rigorous security guarantees which may be attained by incorporating other mechanisms (such as nonces or real-time photographic validation).

Our protocol, if deployed in this case, already incorporates security properties of a basic location proof and has two important advantages: the link fingerprint protects against man-in-the-middle attacks, and the session record contains a record of the data exchanged. This serves as a type of *transaction proof*, including strong guarantees on location, time, direct communication, and data exchanged during the session.

In the next section, we report on experiments to validate our claim that usable link fingerprints can be generated using wireless channel characteristics.

V. EXPERIMENTAL VALIDATION

We used MicaZ motes, running TinyOS, operating in the 2.4 GHz band. These radios provide received signal strength indicator (RSSI) values, a measure of signal power in logarithmic units, which is suitable for generating link fingerprints due to its sensitivity to location, motion, and time. To mimic an actual bodyworn sensor deployment, we mounted the device on a human subject Alice (on the upper right arm as shown in Fig. 3(a)), that communicates with an off-body basestation Bob (pictured in Fig. 3(b)). Our indoor environment is an office space with multiple cubicles, furniture and people. The layout is depicted in Fig. 3(c), marking out the locations of the basestation (Bob) and three eavesdroppers (Eve1, Eve2, Eve3) at a distance of greater than one wavelength away from the legitimate parties. The bodyworn device transmits packets at the rate of 1 packet/second, typical for healthcare devices transmitting physiological information such as heart-rate, ECG, etc. The basestation responds within $10 \sim 20$ milliseconds with an acknowledgement to every message. This routine device communication enables both parties to sample the wireless link in succession and record the RSSI values.

(a) Mobile mote (Alice)   (b) base-station (Bob)

(c) Layout of experimental setup

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

The scheme can be adjusted for different sampling rates as we will discuss later. Furthermore, packet loss is common in mobile environments, and using acknowledgement packets with running counter values has further benefit in that it enables both sensor device and basestation to synchronize their transmissions and accordingly their link fingerprints.

Our experiments consist of two activity modes: *High Activity* where the subject, Alice, walks around the office space to different cubicles, and interacts with other people in the room, and *Low Activity* where she is mostly seated at her cubicle, occasionally getting up to fetch items from nearby cubicles. For each activity, we collect RSSI trace readings for the bodyworn sensor device (Alice), basestation (Bob), and eavesdroppers (Eve1-3), spanning approximately 40 minute periods, which we analyze offline with Matlab.

We provide in Fig. 4 a one-minute sample of the RSSI trace for the *High Activity* experiment. It is observed in Fig. 4(a) that the bodyworn sensor device, Alice, and the basestation, Bob, channel measurements are highly correlated. Eavesdroppers, however, experience a different channel and are unable to replicate the RSSI profile, as shown in Fig. 4(b). As we noted earlier, there is rapid signal decorrelation at distances of over half a wavelength, and independent signals may be assumed for a separation of one or two wavelengths and more. In

the case of the 2.4 GHz band, this indicates that if Eve is at a distance greater than 13 cm of Alice or Bob, she will experience different fading characteristics than the legitimate Alice-Bob channel and not be able to deduce the channel profile. For this reason, research solutions based on wireless link characterization stipulate as part of the threat model that eavesdroppers be situated at least two wavelengths away from the legitimate parties.

To quantify the correlation for channel measurements for the different parties, we compute the Pearson correlation coefficient r:

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

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 is ideal for channel profile characterization in that it measures variations and not absolute values, and is therefore unaffected by offsets in RSSI measurements arising from differences in receiver sensitivities or transmit powers.

The results are presented in Table I. Again, it is observed that there is strong correlation between the legitimate parties (the Alice-Bob correlation is greater than 0.9), whereas it is poor for the eavesdroppers (Alice-Eves correlation is generally below 0.2). We also present results for filtered versions of the

(a) RSSI trace for Basestation (Bob) and Mobile Mote (Alice)

(b) RSSI trace for Basestation (Bob) and Eavesdroppers (Eve)

Fig. 4: Sample of RSSI trace for *High Activity*

TABLE I: Correlation coefficient (r) of RSSI measurements observed by various parties

Experiment	Alice-Bob (r)	Alice-Eve1	Alice-Eve2	Alice-Eve3
High Activity	0.974	0.197	0.088	0.038
Low Activity	0.950	0.129	0.102	0.158
High Activity (filtered)	0.986	0.281	0.118	0.065
Low Activity (filtered)	0.976	0.205	0.152	0.224

channel profile. Filtering is useful because it reduces nonsymmetric discrepancies between the two parties (due to elements such as random noise, sampling delay at the endpoints, etc.) and has been recommended in the literature [23] [32]. For our purposes we use the Savitzky-Golay filter (other filters such as moving average techniques can also be used) and correlation is seen to improve slightly, giving higher confidence (> 0.95 correlation) in the shared fingerprint between Alice and Bob, while still keeping correlation low (< 0.3) for eavesdroppers.

These results are as expected and in good agreement with prior work. We do not provide a thorough characterization of the wireless channel in this paper as it has been extensively documented in the literature: the interested reader is referred to a detailed study in [24] and specifically [33] for the off-body channel.

Considering the strong correlation between the bodyworn sensor device and the basestation, technically the RSSI measurements themselves could be the link fingerprint. Both parties, Alice and Bob, could simply encrypt and sign their RSSI measurements, and a third party, Victor, could compare the RSSI trace results, measure the correlation coefficient (much as we have done), and if he calculates a value $r > 0.9$, he can be certain that the fingerprint is valid.

However, there are issues with using raw RSSI values as link fingerprints: for one, both parties will have to store every RSSI value for every transaction in memory which may not be feasible for memory-constrained sensor devices. For example, the MicaZ motes record RSSI in single byte-sized values and, at a sampling rate of 1 packet/second, would exhaust their 4 KB of RAM in little over an hour. Second, radio usage is a very expensive operation for these devices [43], and these RSSI measurements have to be offloaded from the sensor device as part of the session record, resulting in extra transmission overheads. In the next section we show how the storage, communication and energy overheads can be dramatically reduced by leveraging known quantisation techniques for compressing the RSSI-based link fingerprint.

VI. OPTIMIZATION AND DISCUSSION

Quantization is a digitizing technique that can efficiently distil the raw RSSI data to a much smaller and manageable size. Another advantage is that it has been well-studied in the literature (especially in the context of wireless channel-based secret-key generation [24]) and can be designed to further reduce nonsymmetric noise components in the signals observed by Alice and Bob.

In generating link fingerprints, legitimate communicating parties Alice and Bob sample the wireless channel over

a period of time to gather sufficient channel variation (or entropy) which is then quantized to yield a common bitstring. Quantization mechanisms typically consist of level crossing or ranking techniques and the operator of the scheme can choose one depending on application requirements. We describe an example of each approach here, validate them with our experimental RSSI traces, and compare their properties.

A. Level Crossing Quantization

Fig. 5 depicts a basic level-crossing quantizer (defined in prior work [24]). The bodyworn sensor device and the base-station define an adaptive moving window of size W_Q, within which consecutive (filtered) RSSI readings are processed. For each window, 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 exceeds $q+$, it is encoded as 1, and if less than $q-$, as 0. The thresholds define an exclusion zone and values lying in between them are discarded. This helps to further remove small scale discrepancies between the two endpoints, whereas there is usually very good agreement for excursions larger than the standard deviation. The α parameter allows the operator to adjust quantizer performance to balance between bit generation rate and bit agreement. For our purposes, we use a window size of $W_Q = 5$ and $\alpha = 1$, consistent with prior work.
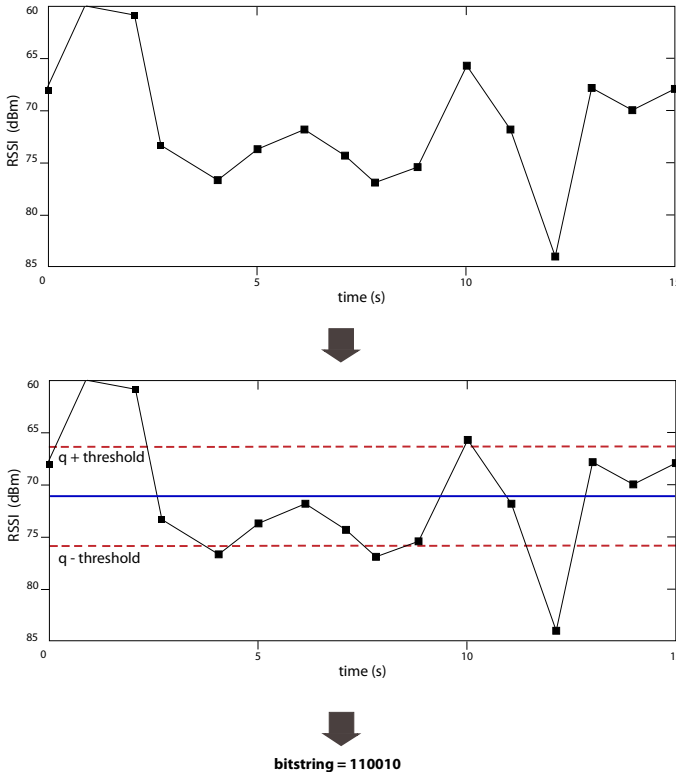
B. Ranking Quantization

A multi-bit ranking quantizer is depicted in Fig. 6. The algorithm sorts the RSSI values in order to divide them into n equal-sized 'buckets' ($n = 4$ in this case). Each RSSI value in the original channel profile can then be encoded with $\log_2 n$ bits . Gray coding is used to number the buckets instead of binary coding, because successive values in Gray coding differ in only one bit, and will therefore limit small RSSI disagreements between Alice and Bob to a single bit per discrepancy.

C. Performance Results

We perform the quantization process to generate fingerprints for the two activity modes for all parties using level crossing and ranking technique, and present results in Table I. We
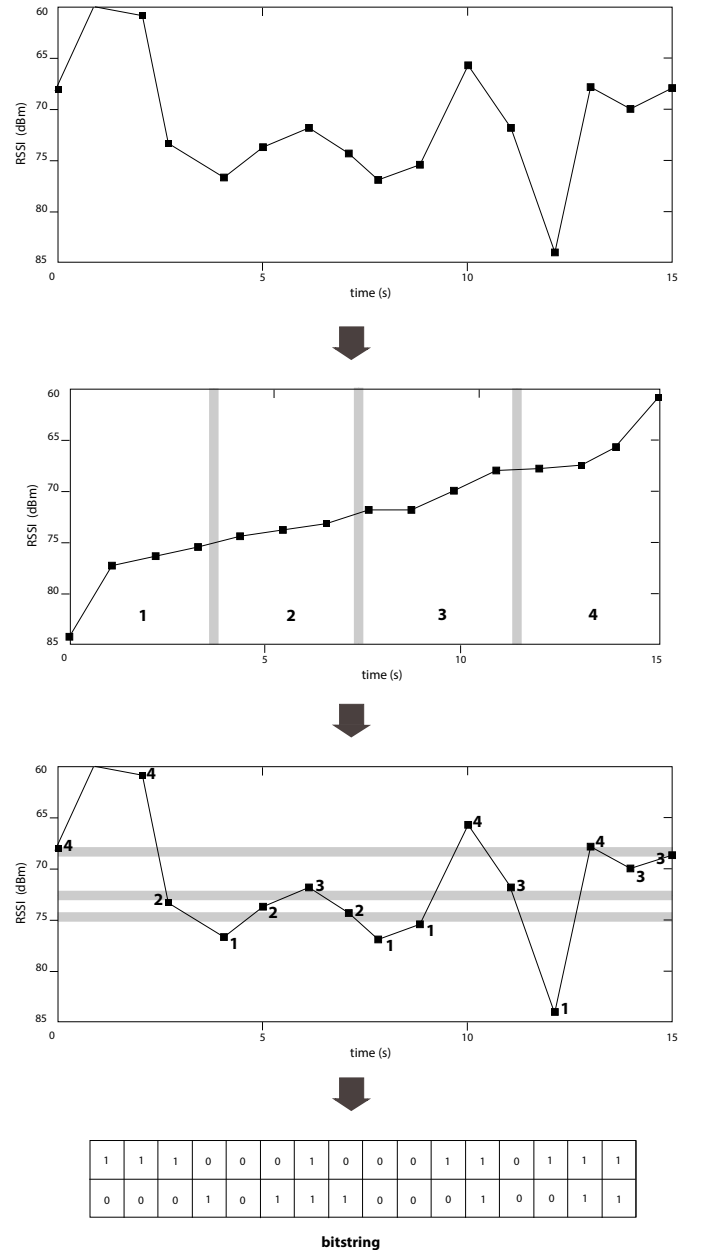

Fig. 5: Level crossing quantization technique

Fig. 6: Ranking quantization technique

Activity (Quantization)	Fingerprint Agreement	Bit (bit/s)	Min. Session Length (mins)	Eve1 Agreement	Eve2 Agreement	Eve3 Agreement	Entropy
High Activity (Level Crossing)	98.40%	0.205	10.41	47.11%	46.48%	47.34%	0.997
Low Activity (Level Crossing)	95.53%	0.139	15.35	46.26%	46.80%	47.60%	0.997
High Activity (Ranking)	93.60%	2	2.13	44.39%	46.92%	48.74%	1
Low Activity (Ranking)	96.08%	2	2.13	50.54%	50.41%	52.92%	1

TABLE II: Link fingerprint performance for experimental scenarios

briefly discuss here our findings and the metrics we use to evaluate our solution:

1) **Bit Agreement**: is the percentage of bits in the fingerprint that are matching between the bodyworn device and the basestation. As can be seen, this is 90% or greater for the legitimate parties and can be used to conclusively validate the link between them. Bit agreement is better in general for the level crossing quantizer because, unlike for the case of ranking where every RSSI value is quantized, the level crossing algorithm discards those values within the exclusion zone that are likely to cause disagreement.

2) **Bit Rate**: is the average number of bits that can be extracted from the channel per unit time. The ranking quantizer performs at a constant rate of 2 bits/s since all of the raw RSSI values are encoded. The level crossing technique exhibits a much lower rate because a single RSSI value can only be encoded to a single bit and also several RSSI values are discarded because they lie in the exclusion zone. There is a tradeoff between bit agreement and rate.

3) **Minimum Session Length**: is the fingerprint length divided by the bit rate. The operator of the scheme can choose the length of the desired fingerprint. For level crossing, depending on the activity mode, it takes approximately 11 to 16 minutes to generate a 128 bit link fingerprint. For the ranking quantizer, which has a much faster rate, a fingerprint can be generated in approximately 2 minutes.

4) **Eavesdropper Agreement**: Fingerprints generated by the eavesdroppers should ideally match with the legitimate parties for 50% of the bits, which we see in our results. This translates to their knowledge of the legitimate fingerprint gained by eavesdropping as being no more useful than a tossing a coin.

5) **Entropy**: is a measure of the inherent randomness or uncertainty in the key. For a random variable *X*, over the set of *n* symbols $x_1, x_2, ..., x_n$, entropy is typically measured as follows:

$$H(X) = -\sum_{i=1}^{n} p(x_i) \log_2 p(x_i) \qquad (2)$$

where $p(x_i)$ is the probability of occurrence of symbol x_i. For binary symbols, a value close to 1 indicates high entropy, which we achieve in our results. Furthermore, the fingerprints we generate clear the entropy tests in NIST test suite [44], a battery of tests with a pass/fail

result, typically used in the literature to confirm randomness of wireless channel-based secrets.

We believe these are encouraging results which validate our proposed approach, and lay the grounds for future work. The tradeoffs between the different quantizers are also highlighted: ranking can be used to build lossless multi-bit quantizers with high bit generation rate, more suited to applications where the average session time between sensor device and basestation is low. A level crossing technique could be used for longer session times, and yields a fingerprint with higher agreement between the two ends. Quantizers could even be customised to prioritize a desired metric as per application requirements.

D. Energy Comparison

In this section, we quantify the benefits of optimization by comparing the energy consumption of our solution running on bodyworn sensor devices with and without quantization. We use energy benchmarks of the Mica motes specified in the literature and estimate the costs of running our protocol. We assume that energy savings on actual bodyworn sensing platforms will be very similar. Relevant benchmark figures are presented in Table. III.

We specifically simulate the scenario described in Section V where the subject, wearing a bodyworn sensor device on his person, works in an indoor office environment. Transactions between the sensor device and basestation are approximately 40 minutes in duration and we use the trace data collected in our experiments to compute energy consumption for both *High* and *Low Activity* modes. Protocol operation on a sensor device consists of symmetric key encryption of the link fingerprint, hashing it as part of the digital signature process, computing the signature, and then transmitting the result.

For the communication cost, researchers [47] measured the energy consumption for transmitting and receiving data at a transmit power of -5 dBm and the effective data rate. This data rate is to be differentiated from the MicaZ claimed data rate (of 250 kbps) and is far less (at 121 kbps), further reduced by the headers and footers appended to the data by the lower layers of the communication stack. For symmetric encryption, we use energy measurements for the hardware

TABLE III: Energy costs for MicaZ motes

Activity	Cost
AES-128 encryption (128 bits) [45]	1.83 μJ
SHA-1 Hash (64 bits) [46]	154 μJ
ECDSA-160 Sign [47]	52 mJ
Transmit 1 bit [47]	0.6 μJ

AES implementation provided by the CC2420 radio on the MicaZ motes. This is a far cheaper and more convenient option than using a software implementation of a symmetric cipher. SHA-1 hashing costs were originally measured for the Mica2 platform which employs the same Atmel Atmega 128L processor as the MicaZ. For the digital signature operation, we use the Elliptic Curve Digital Signature Algorithm (ECDSA) [48] with 160-bit public key size.

We do not factor in energy costs for symmetric key decryption or digital signature verification because the bodyworn sensors are not required to perform these operations. In our protocol solution, only Victor performs fingerprint decryption and verification and furthermore, is not limited by the energy constraints of the sensor device. We also do not add in the cost of optimization, i.e. linear mathematical operations such as filtering, averaging, sorting, computing standard deviation, etc. as these processing costs are orders of magnitude lower than those of the cryptographic operations and data transmission. As an example, the energy required for a compute operation for one clock cycle on the MicaZ mote, is 3.5 nJ as compared to the cost of transmitting one data bit, which is 0.6μJ [47].

As we observe in Table IV, the benefits of optimization are clear: for both *High* and *Low Activity* modes, using raw RSSI values as the link fingerprint requires approximately twice as much energy as when quantization is used. The reason for this is the length of the fingerprint: the larger the fingerprint, the more the costs will add up in terms of hashing and data transmission. AES encryption costs are also proportional to the size of the fingerprint, but since the operation is done in hardware in this case, these costs are orders of magnitude less than that of hashing and transmission.

Level crossing quantization is the most economical mode of fingerprint generation. Computation and transmission costs in this mode for both *High* and *Low Activity* over the forty minute period are minimal, and add less than 1.5 mJ to the ECDSA costs. Ranking, on the other hand, adds up to 15 mJ to the baseline cost.

The ECDSA signing operation is the most expensive and dominates the energy cost. This operation may be considered the baseline cost as it is constant, independent of the size and mode of computation of the fingerprint. It is therefore possible to economize the cost of this operation by compiling the fingerprint over greater periods of time, or distributing it over multiple transactions, or even utilizing signature amortization techniques [14] such as hash chains or Merkle trees.

E. Protocol Enhancements

In this section, we discuss possible enhancements and practical considerations for our protocol.

In some settings, such as hospitals or gyms, where there are multiple basestations and possibilities for **roaming**, a bodyworn sensor device may form link associations with different basestations over a period of time. The sensor device worn by a hospital patient may communicate for the most part with the basestation in the ward, except when the patient visits the hospital cafeteria where it associates with another basestation. Associations may be very brief and frequently disrupted. In this case, the sensor device and individual basestations could maintain running counter values and incremental fingerprints for communications, such that a complete fingerprint may be generated and signed over multiple sessions between the sensor and basestation pair, only when sufficient data has been communicated between the two.

We also note that it is possible to extend this concept to networks with **multiple hops** (such as mobile sensor networks, delay tolerant networks, etc.) and document the entire wireless path for a data item. If the sensor device transmits data to the basestation which in turn forwards it to another device using the wireless channel, each party in the path could generate the associated signature records and maintain receipts. The verifying party, Victor, could map out the entire wireless path by performing the fingerprint verification process for every link, and, in a loose sense, may even be able to track the mobile parties. Mechanisms could even be developed to ensure that malicious parties in the path do not collude with each other.

In case of known environments, such as hospitals and homes, where patterns for packet loss and patient mobility are broadly understood, it may be possible to choose quantizer parameters in advance for optimal performance. However, this could even be an **adaptive** process. Once the RSSI values have been collected by both parties, quantizer parameters (such as the α parameter and number of buckets, n) may be adjusted on the fly to generate a fingerprint of a desired length. However, more experimentation needs to be done to determine the bounds within which quantizer paramters may be modified without compromising the security and integrity of the fingerprint.

It must be pointed out here that the fingerprint is derived from variation in wireless channel characteristics. If both Alice and Bob are stationary and communicating in a **static environment**, it is possible that the fluctuation in the channel profile may be too low to yield an adequate fingerprint. Furthermore, research indicates that it is possible for an attacker to predict certain characteristics of the channel profile by manipulating aspects of the environment around Alice and Bob if they are stationary [24]. As an example, the attacker could periodically block the line-of-sight communication path between Alice and Bob in a careful controlled manner, thereby causing predictable spikes in the channel profile, allowing the attacker to guess certain bits of the fingerprint.

In such scenarios, the communicating parties can still preserve the secrecy and integrity of the fingerprint by using different fingerprint generating techniques which are not predictably reliant on motion, such as employing frequency hopping [49] or UWB communication [50]. These techniques can also be used at very high millisecond sampling rates to very rapidly generate fingerprints for **high speed applications**. One example is that of access control to buildings, where a person can only afford to spend a few seconds waving his access card in front of a card reader to gain access.

We intend to explore these ideas in future work. We are also working on prototyping our solution to study performance across a wider range of environments and activities.

TABLE IV: Comparison of energy costs for quantization mechanisms

Activity (Quantization Mode)	Size (Bytes)	Transmission Costs (mJ)	AES Encrypt (mJ)	SHA-1 Hash (mJ)	ECDSA Sign (mJ)	Total (mJ)
High Activity (Raw RSSI Values)	2392	11.481	0.274	46.046	52	109.801
High Activity (Ranking)	598	2.870	0.068	11.512	52	66.450
High Activity (Level Crossing)	54	0.259	0.006	1.040	52	53.305
Low Activity (Raw RSSI Values)	2270	10.896	0.260	43.698	52	106.853
Low Activity (Ranking)	568	2.726	0.0650	10.934	52	65.725
Low Activity (Level Crossing)	24	0.115	0.003	0.462	52	52.580

VII. CONCLUSION

In this paper we have proposed a data provenance protocol for bodyworn devices that exploits symmetric spatio-temporal characteristics of the wireless channel. Our solution generates unique link fingerprints that we use to form data to wireless link associations. In contrast to existing provenance mechanisms which operate on a per packet basis, this solution generates provenance on a per session basis, which minimizes the use of cryptographic techniques and associated overheads. The link fingerprints can be built using routine data transmissions, they are unique to the two communicating parties and cannot be deduced in detail by an eavesdropper situated at a distance. Our provenance solution also provides system-wide accountability and non-repudiation.

We performed experiments using bodyworn devices in an indoor office space to demonstrate the high correlation in channel measurements between the two endpoints. We suggest two optimization techniques, level crossing and ranking, to quantize the raw RSSI values to a manageable size, we highlight the benefits of each in terms of energy consumption, and we discuss possibilities for adapting the fingerprinting process for different application requirements. We believe this is a promising first step in using wireless link-based techniques to secure data provenance.

VIII. ACKNOWLEDGEMENTS

REFERENCES

[1] Nike+ FuelBand. Retrieved on 21 July, 2013. http://www.nike.com/us/en_us/c/nikeplus-fuelband.
[2] Fitbit Flex. Retrieved on 21 July, 2013. http://allthingsd.com/20130715/fitbit-flex-vs-jawbone-up-and-more-a-wearables-comparison/.
[3] Toumaz Technology Ltd. *Sensium Life Platform*. http://www.toumaz.com/page.php?page=sensium_intro.
[4] David Szondy. Bracelet Uses Social Network to Protect Civil Rights Activists. 7 April, 2013. http://goo.gl/MIEk2.
[5] Waiting for Apple's iWatch. 22 March, 2013. http://goo.gl/xTsQO.
[6] Disposable Wireless Sensor Market Shows Signs of Life - Healthcare Shipments to Reach 5 Million in 2018. 3 March, 2013. http://goo.gl/V9QoP0.
[7] A. S. Wander, N. Gura, H. Eberle, V. Gupta, and S. C. Shantz. Energy Analysis of Public-key Cryptography for Wireless Sensor Networks. In *3rd IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 324–328, 2005.
[8] U. Braun, A. Shinnar, and M. I. Seltzer. Securing Provenance. In *USENIX Summit on Hot Topics in Security (HotSec)*, 2008.
[9] D. J. Pohly, S. McLaughlin, P. McDaniel, and K. Butler. Hi-Fi: Collecting High-fidelity Whole-system Provenance. In *28th ACM Annual Computer Security Applications Conference*, pages 259–268, 2012.
[10] B. Przydatek, D. Song, and A. Perrig. SIA: Secure Information Aggregation in Sensor Networks. In *1st ACM International Conference on Embedded Networked Sensor Systems*, pages 255–265, 2003.

[11] H. Lim, Y. Moon, and E. Bertino. Provenance-based Trustworthiness Assessment in Sensor Networks. In *Seventh ACM International Workshop on Data Management for Sensor Networks*, pages 2–7, 2010.
[12] H. Lim, G. Ghinita, E. Bertino, and M. Kantarcioglu. A Game-Theoretic Approach for High-Assurance of Data Trustworthiness in Sensor Networks. In *28th IEEE International Conference on Data Engineering (ICDE)*, pages 1192–1203, 2012.
[13] M. A. Chowdhury, W. Mciver, and J. Light. Data Association in Remote Health Monitoring Systems. *IEEE Communications Magazine*, 50(6):144–149, 2012.
[14] S. T. Ali, V. Sivaraman, and D. Ostry. Authentication of Lossy Data in Body-sensor Networks for Healthcare Monitoring. In *9th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, pages 470–478, 2012.
[15] J. C. Sriram, M. Shin, T. Choudhury, and D. Kotz. Activity-aware ECG-based Patient Authentication for Remote Health Monitoring. In *ACM International Conference on Multimodal Interfaces*, pages 297–304, 2009.
[16] K. B. Rasmussen and S. Capkun. Implications of Radio Fingerprinting on the Security of Sensor Networks. In *Third International Conference on Security and Privacy in Communications Networks, SecureComm*, pages 331–340. IEEE, 2007.
[17] B. Shebaro, S. Sultana, S. R. Gopavaram, and E. Bertino. Demonstrating a Lightweight Data Provenance for Sensor Networks. In *ACM Conference on Computer and Communications Security*, pages 1022–1024, 2012.
[18] M. G. Zapata. Secure Ad hoc On-demand Distance Vector Routing. *ACM SIGMOBILE Mob. Comput. Commun. Rev.*, 6(3):106–107, 2002.
[19] W. C. Jakes. *Microwave Mobile Communications*. Wiley, 1974.
[20] Z. Li, W. Xu, R. Miller, and W. Trappe. Securing Wireless Systems via Lower Layer Enforcements. In *5th ACM Workshop on Wireless Security*, pages 33–42, 2006.
[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] S. N. Premnath, P. L. Gowda, S. K. Kasera, N. Patwari, and R. Ricci. Secret Key Extraction using Bluetooth Wireless Signal Strength Measurements. In *IEEE International Conference on Sensing, Communications and Networking (SECON)*, 2014.
[23] S. Mathur, W. Trappe, N. Mandayam, C. Ye, and A. Reznik. Radio-telepathy: Extracting a Secret Key from an Unauthenticated Wireless Channel. In *14th ACM International Conference on Mobile Computing and Networking*, pages 128–139, 2008.
[24] S. Jana, S. N. Premnath, M. Clark, S. K. Kasera, N. Patwari, and S. V. Krishnamurthy. On the Effectiveness of Secret Key Extraction from Wireless Signal Strength in Real Environments. In *15th ACM Annual International Conference on Mobile Computing and Networking*, pages 321–332, 2009.
[25] 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.
[26] S. L. Cotton and W. G. Scanlon. An Experimental Investigation into the Influence of User State and Environment on Fading Characteristics in Wireless Body Area Networks at 2.45 GHz. *IEEE Transactions on Wireless Communications*, 8(1):6–12, 2009.
[27] 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.
[28] 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.
[29] L. Shi, M. Li, S. Yu, and J. Yuan. Bana: Body Area Network Authentication Exploiting Channel Characteristics. In *5th ACM conference on Security and Privacy in Wireless and Mobile Networks*, pages 27–38, 2012.

[30] L. Shi, J. Yuan, S. Yu, and M. Li. ASK-BAN: Authenticated Secret Key Extraction Utilizing Channel Characteristics for Body Area Networks. In *ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2013.

[31] 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*, 2010.

[32] L. Yao, S. T. Ali, V. Sivaraman, and D. Ostry. Improving Secret Key Generation Performance for On-body Devices. In *6th International Conference on Body Area Networks*, pages 19–22. ICST, 2011.

[33] S. T. Ali, V. Sivaraman, and D. Ostry. Zero Reconciliation Secret Key Generation for Body-worn Health Monitoring Devices. In *5th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pages 39–50, 2012.

[34] J. Tang and P. Fan. A RSSI-based Cooperative Anomaly Detection Scheme for Wireless Sensor Networks. In *International Conference on Wireless Communications, WiCom*, pages 2783–2786. IEEE, 2007.

[35] N. Patwari and S. K. Kasera. Robust Location Distinction using Temporal Link Signatures. In *13th Annual ACM International Conference on Mobile Computing and Networking*, pages 111–122, 2007.

[36] S. Mathur, R. Miller, A. Varshavsky, W. Trappe, and N. Mandayam. ProxiMate: Proximity-based Secure Pairing using Ambient Wireless Signals. In *ACM International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2011.

[37] A. Kalamandeen, A. Scannell, E. de Lara, A. Sheth, and A. LaMarca. Ensemble: Cooperative Proximity-based Authentication. In *8th Annual ACM International conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 331–344, 2010.

[38] Y. Chen, J. Yang, W. Trappe, and R. P. Martin. Detecting and Localizing Identity-based Attacks in Wireless and Sensor Networks. *IEEE Transactions on Vehicular Technology*, 59(5):2418–2434, 2010.

[39] J. Yang, Y. Chen, W. Trappe, and J. Cheng. Detection and Localization of Multiple Spoofing Attackers in Wireless Networks. *IEEE Transactions on Parallel and Distributed Systems*, 24(1):44–58, 2013.

[40] S. Brands and D. Chaum. Distance-bounding Protocols. In *Advances in Cryptology (EUROCRYPT'93)*, pages 344–359. Springer, 1994.

[41] K. B. Rasmussen and S. Capkun. Realization of RF Distance Bounding. In *USENIX Security Symposium*, pages 389–402, 2010.

[42] S. Saroiu and A. Wolman. Enabling New Mobile Applications with Location Proofs. In *10th workshop on Mobile Computing Systems and Applications*, page 3. ACM, 2009.

[43] D. Culler, D. Estrin, and M. Srivastava. Guest editors' introduction: Overview of sensor networks. *Computer*, 37:41–49, 2004.

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

[45] J. Lee, K. Kapitanova, and S. H. Son. The Price of Security in Wireless Sensor Networks. *Computer Networks*, 54(17):2967–2978, 2010.

[46] C. Chang, S. Muftic, and D. J. Nagel. Measurement of Energy Costs of Security in Wireless Sensor Nodes. In *16th International Conference on Computer Communications and Networks, ICCCN*, pages 95–102. IEEE, 2007.

[47] G. D. Meulenaer, F. Gosset, F. Standaert, and O. Pereira. On the Energy Cost of Communication and Cryptography in Wireless Sensor Networks. In *International Conference on Wireless and Mobile Computing, Networking and Communications, WIMOB'08.*, pages 580–585. IEEE, 2008.

[48] Darrel Hankerson, Scott Vanstone, and Alfred J Menezes. *Guide to elliptic curve cryptography*. Springer, 2004.

[49] M. Wilhelm, I. Martinovic, and J. B. Schmitt. Secure Key Generation in Sensor Networks Based on Frequency-Selective Channels. *IEEE Journal on Selected Areas in Communications*, 31(9):1779–1790, 2013.

[50] 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):364–375, 2007.

Syed Taha Ali Syed Taha Ali did his BSc. (Eng) from GIK Institute of Engineering Sciences and Technology, Pakistan in 2002. He did his MS and PhD in Electrical Engineering from the University of New South Wales, Australia, in 2006 and 2012 where he currently works with the Network Research Group in the School of Computer Science and Engineering. His research interests include body area networks, software defined networking, and network security.

Vijay Sivaraman 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 Assoc. Professor at the University of New South Wales, Australia. His research interests include software defined networking, optical networking, packet switching and routing, and wireless sensor networks.

Diethelm Ostry 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, optical packet networks and security in body sensor networks.

Gene Tsudik Gene Tsudik is a "Lois and Peter Griffin" Professor of Computer Science at the University of California, Irvine (UCI). He obtained his PhD in Computer Science from USC in 1991 for research on firewalls and Internet access control. Before coming to UCI in 2000, he was a Project Leader at IBM Zurich Research Laboratory (1991-1996) and USC Information Science Institute (1996-2000). Over the years, his research interests included: routing, firewalls, authentication, mobile networks, secure e-commerce, anonymity ad privacy, group communication, digital signatures, key management, mobile ad hoc networks, as well as database privacy and secure storage. He currently serves as Director of Secure Computing and Networking Center (SCONCE) and Vice-Chair of the Computer Science Department. In 2007, he was on sabbatical at the University of Rome as a Fulbright Senior Scholar. Since 2009, he is the Editor-in-Chief of ACM Transactions on Information and Systems Security (TISSEC).

Sanjay Jha Sanjay K. Jha is Professor and Head of the Network Group at the School of Computer Science and Engineering at the University of New South Wales. His research activities cover a wide range of topics in networking including Wireless Sensor Networks, Adhoc/Community wireless networks, Resilience and Multicasting in IP Networks and Security protocols for wired/wireless networks. Sanjay has published over 160 articles in high quality journals and conferences. He is the principal author of the book Engineering Internet QoS and a co-editor of the book Wireless Sensor Networks: A Systems Perspective. He served as an associate editor of the IEEE Transactions on Mobile Computing (TMC) and he currently serves on the editorial board of the ACM Computer Communication Review (CCR)