OPEN ACCESS

University of BRISTOL

Peer reviewed version

Link to published version (if available):
10.1109/WF-IoT.2015.7389068

Link to publication record in Explore Bristol Research
PDF-document

University of Bristol - Explore Bristol Research
General rights

Mitigating Packet Loss in Connectionless Bluetooth Low Energy

Evgeny Tsimbalo, Xenofon Fafoutis, Evangelos Mellios, Mo Haghighi,
Bo Tan, Geoffrey Hilton, Robert Piechocki and Ian Craddock
Department of Electrical and Electronic Engineering
University of Bristol
{e.tsimbalo, xenofon.fafoutis, evangelos.mellios, mo.haghighi, b.tan, geoff.hilton, r.j.piechocki, ian.craddock}@bristol.ac.uk

Abstract—In this paper, we investigate various methods to combat packet loss in a residential communication system based on the Bluetooth Low Energy (BLE) standard, focusing on BLE's connectionless mode (undirected advertising) in which no retransmissions are possible. We start by introducing two orthogonally polarised antennas at the receiver, thus improving the probability of successful reception. This is followed by enabling error correction using redundancy introduced by the Cyclic Redundancy Check (CRC) code of BLE. The CRC error correction is based on a novel approach of applying iterative decoding algorithms. We then consider a BLE system deployed in a residential environment and utilise the presence of multiple receivers that are necessary to provide coverage. These three techniques come at no cost for the transmitter, thus preserving its energy efficiency. The final technique deals with error control coding in the application layer, in which some redundancy is added at the transmitter before data is sent to the physical layer. By combining all four methods, a distributed error correction algorithm is developed. Using real BLE packets collected in a typical 2-storey house, it is shown that the designed system can correct 80% of all corrupted packets.

Keywords-Bluetooth Low Energy; error correction; packet loss

I. Introduction

Smart homes and wearable technologies are widely considered as a possible answer to the challenges of national health systems. Indeed, the aging population and the rise of chronic conditions, such as depression or diabetes, drive the cost of healthcare to unsustainable levels. In this context, advances in wearable computers, Wireless Sensor Networks (WSNs) and Ambient Assisted Living (AAL), provide means that support self-management of health and well-being, and facilitate timely interventions. A key challenge of developing long-term services for such target groups is energy consumption. Although acceptable in fashionable gadgets, healthcare technologies cannot depend on the user for regular maintenance, such as battery charging.

Energy-efficient communication is generally considered a key to a long battery lifetime. Bluetooth Low Energy (BLE) [1] is a compelling choice for such applications, primarily due to its very low power nature and its widespread integration. BLE defines two modes of operation. The first is based on a connectionless unidirectional communication scheme where the transmitter is broadcasting to a receiver. The second implements the traditional connection-oriented Bluetooth scheme. The simplicity of the connectionless BLE makes it an ideal choice for long-term monitoring in a residential environment. Its advantages can be summarised as follows. Firstly, it conveys data through three advertisement channels that are carefully designed not to interfere with the popular IEEE 802.11 channels 1, 6 and 11 [2][3]. Secondly, when the user is moving, the handovers to other receivers are seamless, quick and without overhead. Lastly, its minimal design introduces low overhead and is energy-efficient, due to absence of connection establishment and channel listening. On the negative side, the key weakness of connectionless BLE is the lack of acknowledgements and retransmissions. Channel errors are experienced as application-layer data loss.

In this paper, we investigate techniques that mitigate packet loss in connectionless BLE. We consider a broadcasting network, where an energy-constrained wearable computer streams data to a constraint-free smart infrastructure of multiple receivers. Instead of energy-consuming retransmissions, the key idea of the proposed system is to maintain the wearable sensor simple and energy-efficient, by applying mechanisms that mitigate packet loss in the infrastructure. In particular, we investigate four different techniques, namely (i) orthogonally polarised antennas, (ii) CRC (Cyclic Redundancy Check) error correction, (iii) application-layer coding, and (iv) multiple spatially-distributed receivers. The contribution of this work lies in its unique combination of all four techniques into a distributed system that is suitable for deployment in residential environments. Experiments in a residential environment show that the designed system mitigates packet loss and significantly improves the reliability of connectionless BLE. To the extent of our knowledge, this work documents, for the first time in the literature, the effectiveness of the aforementioned techniques in a practical BLE broadcasting network and their benefits in comparison to off-the-shelf BLE.

The remainder of the paper is structured as follows. Section II presents the four applied techniques. Section III evaluates each of the methods separately. Section IV combines the techniques in a full system design and demonstrates the full-system performance. Section V concludes the paper.

II. Packet Loss Mitigating Techniques

In a normal BLE-based communication system, data bits from the upper layers are first encapsulated into physical

(PHY) layer packets. The CRC encoder adds additional 24 redundant bits to each packet that are used to detect errors at the receiver (CRC-24). The encoded data is then modulated and up-converted onto a carrier frequency. At the receiver, the arrived signal is down-converted and demodulated. For each packet, it is then checked whether the packet contains bit errors. If there is at least one bit error, the packet is discarded. Otherwise, the packet is forwarded to the upper layers of the receiver. In the connectionless mode, dropped packets are experienced as application-layer data loss. It is clear that an off-the-shelf connectionless BLE system performs well only if the probability of bit error is very low.

A. Orthogonally Polarised Antennas

One of the simplest ways to improve the reliability of a practical BLE system is to introduce two orthogonally polarised antennas at the receiver [4]. In practice, this can realised by incorporating two parallel receiving radios, instead of just one, such that one antenna is polarised vertically and another horizontally. Considering that one end of the link, i.e. the wearable sensor, is unpredictably polarised, as it is following the movements of the user, such antenna arrangement would increase the probability of successful reception. The input streams from the two radios are required to be combined into one stream after the CRC check. The cost of the solution is a minimal increase in computation complexity at the receiver. However, there are no additional costs for the energy-constrained transmitter.

B. CRC Error Correction

While the CRC code employed by BLE is used for error detection only, it has an inherent error correction potential due to redundancy it introduces to transmitted data. CRC error correction would directly improve the packet reception rate at the receiver, thus, for example, allowing it to be placed further from the transmitter.

In our previous works [5][6], we introduced a novel approach to the error correction of CRC codes based on iterative decoding techniques. In particular, it was shown how the BLE CRC code can be converted to an equivalent Low Density Parity Check (LDPC) code by making its parity check matrix sparse using the algorithm proposed in [7]. The most widely used decoding approach for LDPC codes is based on Belief Propagation (BP) [8]. As an alternative to BP, a linear programming approach to the decoding of LDPC codes was proposed in [9]. The resulting algorithm known as the Alternating Direction Method of Multipliers (ADMM) was developed in [10] and [11]. In [5][6], we studied both BP- and ADMM-based decoding algorithms in the context of the BLE CRC code. The ADMM-based approach showed a superior performance to that of BP, so it is employed in this paper.

The iterative decoding methods described above use soft information about the bits to be decoded in the form of log-likelihood ratios (LLRs). The LLR of a bit indicates how close the bit value is to 1 or 0 based on channel measurements. In a practical BLE system, the LLRs are not directly available at the receiver, since the demodulator produces hard decisions in the form of 1s and 0s. In [5], it is shown that the transmitter, channel and receiver can be modelled as an equivalent binary symmetric channel (BSC) that flips some of the bits encoded by the CRC code with a probability p. The LLR for the i-th received bit r_i can be calculated as

$$\gamma_i = (2r_i - 1)\ln\left(\frac{1-p}{p}\right). \tag{1}$$

In a practical BLE system, some indication of the quality of the received signal is usually available. For example, in our platform [12], the Received Signal Strength Indication (RSSI) is available for every received packet. Therefore, a statistics on the bit flipping probability (or the bit error rate, BER) p as a function of the RSSI level can be measured in advance by sending some known packets. This statistics can then be used as a loop-up table to calculate LLRs in a real scenario.

C. Application-Layer Coding

The packet loss mitigation techniques considered so far did not change the way the information is transmitted. At the same time, it is clear that if some error correction coding was introduced among the data bits prior to transmission, it would allow more packets to be corrected than by means of CRC correction only. We will call this approach Application-Layer Coding (ALC), since it applies to the actual application data bits, before sending them to the PHY layer.

For a general communication system, error correction coding offers a compromise between the amount of data that can be corrected and the level of redundancy. In one extreme case, no redundancy offers the maximum bandwidth utilisation, but at the same time the highest error rate. In another extreme case, data can be transmitted repetitively (using so called repetition coding), which provides excellent error correction capabilities, but the bandwidth utilisation reciprocal to the repetition factor.

Compared with a traditional, non-constrained communication system, low energy systems have an additional factor to take into account - the energy consumed to transmit data. It is therefore natural to define the energy efficiency of a system as a ratio between the energy available at the transmitter and the amount of information reached the receiver. If error correction coding is applied, the encoding process will consume some additional energy. Let n be the total number of bits after encoding. The number of original message bits can be expressed as $m = nr$, where $r \leq 1$ is the code rate. The energy efficiency can now be defined as

$$\eta = \frac{E + \delta}{m(1 - MER)} = \frac{E + \delta}{nr(1 - MER)}, \tag{2}$$

where E is the energy required to transmit an uncoded message, δ is additional energy consumed by the encoding process and MER is the message error rate after error correction. By fixing E and n, (2) can be used to benchmark and compare the performance of different coding techniques: the lower the η, the more efficient the technique is.

In this paper, we exploit the fact that the transmitter sends three advertisement packets on different carrier frequencies,

Fig. 1. The wearable sensor prototype (left) and the receiver that employs two orthogonally polarised antennas (right).

which can be viewed as three independent channels. The natural options are therefore to encapsulate *unique* message bits into each packet (no ALC, $r = 1$), to encapsulate the *same* message bits into each packet (repetition coding, $r = 1/3$) or to employ some encoding between messages bits before sending them to the PHY layer. For the latter case, we propose a simple encoding scheme based on the XOR operation: a message is split into two parts, A and B, with A being sent on the first channel, B on the second, and $A \oplus B$ on the third, where $\oplus$ denotes a bit-wise XOR operation. The resulting code rate is $r = 2/3$. At the receiver, it is clear that to receive the message correctly, it is sufficient to successfully receive any two out of the three packets. Let p denote the PHY packet error rate (PER). Assuming that each packet experiences an independent loss, the MER for the XOR-based ALC can be calculated as follows:

$$MER_{XOR} = p^3 + 3(1-p)p^2 = 3p^2 - 2p^3. \quad (3)$$

The MER for repetition coding and the uncoded systems can be similarly obtained as follows:

$$MER_{rep} = p^3; \quad (4)$$

$$MER_{unc} = 1 - (1-p)^3 = 3p - 3p^2 + p^3. \quad (5)$$

By substituting (3), (4) and (5) to (2) and using the corresponding code rates, the three techniques can be compared in terms of the energy efficiency.

D. Multiple Receivers

In this work we consider a case when a wearable broadcasts to several receivers, which in turn forward traffic to a central device, or a hub. Since the transmission power of BLE radios is typically low, multiple receivers are necessary to provide signal coverage over the area in question, for example a residential house. Leveraging the broadcasting nature of connectionless BLE, traffic combining can be performed in the central hub. Indeed, not only the coverage is increased, but also different pieces of a broken message (if ALC is enabled) may reach different receivers. This would increase the probability of successful reception and allow the wearable to operate at smaller transmission power levels compared with a single receiver scenario.

III. Performance Evaluation

In this section, we experimentally evaluate the considered packet loss mitigation techniques. The experiments are conducted using the prototypes of the SPHERE (a Sensor Platform
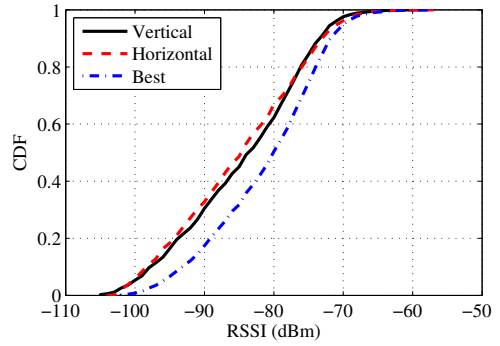

Fig. 2. CDF of the RSSI for different antenna arrangements.

for Healthcare in a Residential Environment) infrastructure [12] that is based on BLE for body-centric communications. The transmitter is the prototype wearable sensor shown in Fig. 1 (left). We refer the reader to [12] and [13] for details on the hardware and employed patch antenna respectively. Using connectionless BLE, the wearable sensor broadcasts data at 5 Hz in all three advertisement channels. We use the maximum allowed packet size (39 bytes). The receiver, incorporates two radios that employ two orthogonally polarised dipole antennas, working in parallel, as shown in Fig. 1 (right). For details on the receiver antennas we refer the reader to [14].

Unless otherwise noted, the presented evaluations are based on data collected in a 2-storey house in the city of Bristol, UK. For the remainder of this paper, we will refer to these data sets as the *SPHERE house measurements*. For these measurements, three receiver units are deployed in the house, programmed to log the RSSI, the outcome of the CRC check, and the hex code of corrupted received packets.

A. Orthogonally Polarised Antennas

The first set of experiments is designed to evaluate the gain from using two antennas with orthogonal polarisations at the receiver. The evaluation uses part of the SPHERE house measurements, in which a user, with the wearable sensor mounted on his wrist, performs random walks within each of the 7 rooms of the house for 70 minutes (10 minutes in each area). The transmission power of the wearable sensor is set at -20 dBm.

Fig. 2 shows the Cumulative Distribution Function (CDF) of the RSSI of the received packets that correspond to horizontally and vertically polarised antenna respectively. As the user moves randomly within house, a different receiver antenna aligns to the polarisation of the transmitted signal. Using two orthogonally polarised antennas, our system keeps the packet that is received with the highest received signal on a per-packet basis. Fig. 2 also plots the CDF of the RSSI after applying two orthogonally polarised antennas. Without any change in the energy-constrained wearable sensor, this technique accounts for an approximately 3 dB median improvement in the RSSI.
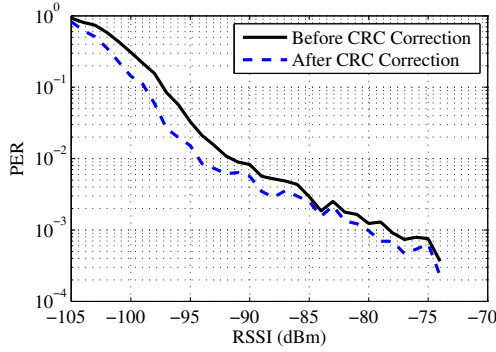
Fig. 3. Packet Error Rate (PER) for various RSSI levels before and after CRC correction.

Fig. 4. Energy consumption per bit for different transmission power levels, in cases of no redundancy, repetition coding and XOR coding ($p = 0$).

B. CRC Error Correction

The evaluation of the CRC Error Correction is based on a larger set of corrupted packets that is collected in an office environment. In this case, a receiver was fixed on a desk, while the wearable transmitter was positioned in different locations within the office and the surrounding areas. To minimise interference, the data collection sessions were performed at off-peak hours when the office was empty. In a period of over 10 hours, we collected a total of approximately $400,000$ packets, $6,000$ of which were corrupted.

Fig. 3 shows the Packet Error Rate (PER) for each RSSI level before any correction algorithm is applied (solid line). We can observe that without error correction, the 1% PER threshold is at -91 dBm. On the other hand, when ADMM is enabled, a gain of up to 3 dB can be achieved, i.e. the 1% PER threshold is at -94 dBm. In terms of the PER, error correction significantly improves the reliability, correcting up to 60% of corrupted packets at -97 dBm. Again, such improvements come at no additional cost at the transmitter.

C. Application-Layer Coding

In contrast to the previously investigated techniques, coding does affect the energy consumption of the transmitter. Considering the maximum permitted packet size, i.e. 39 bytes, each of the three BLE advertisement has a payload of 24 bytes. Therefore a triple advertisement carries 72 payload bytes ($n = 576$). The current profile of a triple advertisement event is estimated by measuring the voltage drop across a $10\ \Omega$ resistor in series with the positive side of the supply, considering all available transmission power settings. The energy consumption of a triple advertisement event (E) is then derived by estimating the integral of the current profile over time and multiplying it by the supply voltage. For instance, at 4 dBm, the energy consumption is estimated at $62\ \mu$J. In the case of XOR-based ALC, an additional $0.3\ \mu$J was recorded to be consumed by the encoding process. Assuming an error-free channel ($p = 0$) and based on (2), Fig. 4 shows the energy consumption per bit for the cases of no redundancy, packet repetition and XOR-based ALC, for various transmission power settings.

Assuming the maximum available transmission power setting (4 dBm), we next consider channel errors. In particular,
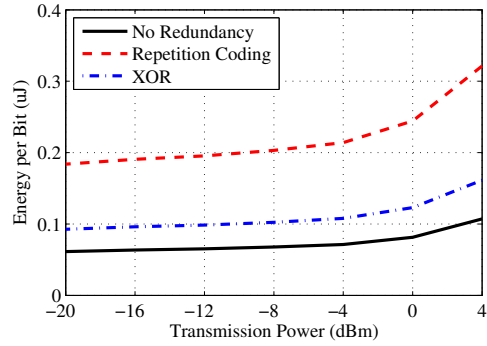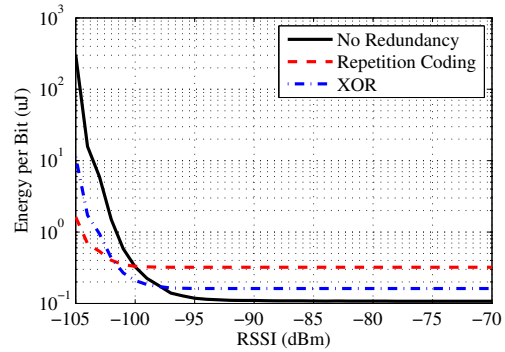

Fig. 5. Energy consumption per bit for different RSSI levels, in cases of no redundancy, repetition coding and XOR coding (4 dBm transmission power).

we use the PER measured in Section III-B (without CRC error correction) to calculate the energy efficiency per bit (2) based the MER of all three schemes given in (3), (4) and (5). Fig. 5 plots the results as a function of the RSSI. It can be observed that each method performs best in a separate RSSI range. As expected, at high signal levels, any coding is redundant, resulting in some increase in the energy consumption. Repetition coding has the highest energy efficiency only when the signal level is extremely low, due to its superior error correction capability. The XOR-based method provides the lowest energy per bit between -102 and -97 dBm, while being more energy efficient than the repetition code at high signal levels. We note that the optimum choice and long-term energy savings in a real scenario will depend on a distribution of the received signal level. However, without *a priori* knowledge of such statistics, coding methods similar to XOR can be a perfect solution.

D. Multiple Receivers

Finally, we quantify the benefit of multiple receivers. The evaluation uses part of the SPHERE house measurements, which include 12 links, classified based on the number of the walls between the receiver unit and the user with the wearable sensor mounted on his wrist. Fig. 6 shows the CDFs of the RSSI of each one of these classes. The transmission power of the wearable sensor is set at 4 dBm. Yet, the wireless performance of lower transmission power levels can be derived
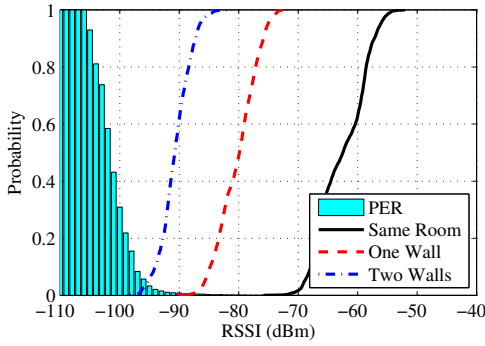
Fig. 6. The CDFs of the RSSI for links that are in the 0, 1 and 2 walls from the receiver show how multiple observers can provide full-house coverage.

by shifting to the left the corresponding CDF. It can be observed that, in the median case, the RSSI of the adjacent-room links is 15 dB lower than the same-room links. A second wall accounts for an additional 10 dB attenuation.

The same figure also plots the PER measured in Section III-B as a bar chart. The results quantify the fact that the bigger a residential environment is, the more receiver units are required for full-house coverage. Combining the results of Fig. 4 and Fig. 6, we can also quantify the trade-off between wireless coverage and energy-efficiency. Deploying additional observers does not increase the energy consumption at the wearable transmitter. Therefore, energy consumption is minimised by setting the transmission power to the lowest level and using multiple observers, along with the other presented techniques, to mitigate packet loss.

IV. System Design and Performance

In this section, we combine all error correction techniques considered so far and implement them in a distributed system consisting of multiple receivers, or access points (AP), and a central hub. To this end, we establish two performance goals. Primarily, we aim at minimising the PER, thus maximising the reliability of the system. Secondly, we spread the computational burden across the network and reduce the traffic between the devices.

A. System design

Combining packets from two antennas with orthogonal polarisations, due to its simplicity, can be implemented on each AP. The combining process can be described as follows. If the same packet is received in both polarisation and if at least one copy is correct (i.e. it passed the CRC check), the other copy is discarded. In this way, the amount of forwarded traffic is reduced. If both copies are incorrect, they are both forwarded further, so that the CRC error correction algorithm can combine the reliability information. Combining packets from multiple APs in the central hub can be implemented in the same manner.

Of all four techniques, CRC error correction is the most computationally intensive. Since the same packet can potentially be received by all APs and in both polarisations,

attempting to correct errors in each copy of the packet would lead to unnecessary processing. On the other hand, combining all copies of the packet and removing unnecessary duplicates in the central device can significantly reduce the total number of corrupted packets. Therefore, it is reasonable to perform CRC error correction block in the central hub. By contrast with the single receiver scenario, soft information about bits should now be computed based on multiple copies of the same packet. Assuming that all received copies are statistically independent, (1) can be rewritten as

$$\gamma_i = \sum_j (2r_{i,j} - 1) \ln \left(\frac{1 - p_j}{p_j} \right), \qquad (6)$$

where $r_{i,j}$ denotes the i-th bit of the j-th copy of the packet and p_j is the RSSI-based bit-flip probability. By summing up the LLRs for statistically independent observations, the reliability estimation is averaged out.

Application-layer (AL) decoding, by definition, is normally applied at the end of a signal processing chain. Following the strategy described in Section II-C, let the encoded sequence of packets be A, B and $A \oplus B$. It can be observed that such decoding procedure, due to its simplicity, could be applied not only after CRC error correction, but also after polarisation combining, with the aim of reducing the amount of traffic sent to the central hub: if any two out of three packets are correct, the third packet becomes redundant. Similarly, after packets from all receivers are combined in the central hub, it may be possible to reconstruct both A and B without resorting to CRC correction. In total, AL decoding can be applied after each of the other three techniques described above.

Algorithm 1 summarises the order and brief description of the error correction techniques.

Algorithm 1 Distributed error correction.

Input: Encoded message comprising of packets A, B and $A \oplus B$ fully or partly received by multiple APs.
Output: Decoded packets A and B.
Local AP processing: on each AP do:
Step 1: **Polarisation combining**: for each unique received packet:
    if only one copy is received, forward it;
    elseif both copies are not correct, forward both;
    else forward the correct copy and discard the other.
Step 2: **AL decoding**: attempt to decode A and B. If successful, forward both and discard redundant copies. If not, forward all copies.
Central device:
Step 3: **Combining from multiple APs**: for each unique received packet:
    if only only one copy is available, forward it;
    elseif all copies are not correct, forward all;
    else forward just one correct copy and discard others.
Step 4: **AL decoding**: repeat Step 2.
Step 5: **CRC correction**: apply Algorithm 1 from [5].
Step 6: **AL decoding**: repeat Step 2.

B. Performance

The performance of the designed system is evaluated next. The evaluation is based on a set of SPHERE house measurements described in Section III-A. In this case, the transmission power of the wearable sensor is set at -4 dBm.
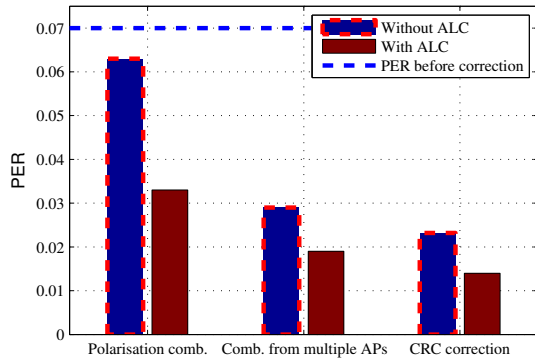
Fig. 7. PER after each stage of correction, with and without ALC.

Fig. 7 illustrates the performance of the system in terms of the PER after each stage of error correction. The performance of the system without ALC is also shown as a reference. The initial PER prior to error correction was recorded at 7% and is shown on the graph as a straight line. It should be noted that the initial PER was calculated over the logged packets only, without taking into account packets that did not reach the APs in the first place. It can be observed that polarisation combining reduces the PER by 10%, which suggests some degree of correlation between the two polarisations. Indeed, in most cases the transmitted signal's polarisation is neither distinctively horizontal nor distinctively vertical. At the same time, it can be seen that combining packets from multiple APs has a dramatic effect on the PER, reducing it further by 54% without ALC and by 42% with ALC. CRC error correction provides an additional reduction of 20% in the PER.

When ALC is enabled and decoding is performed after each of the three other stages, it has a profound effect on the performance and the amount of traffic sent from the APs to the central device. Indeed, the final PER with ALC is only 1.4%, 40% lower than without ALC. In terms of the traffic flow, it was recorded that the number of packets sent to the central device was reduced by 30% when AL decoding was enabled locally at each AP. Finally, the number of packets processed by the CRC error correction block was reduced by 45%, which significantly decreases the computational load in the central hub. Overall, all techniques combined achieve 80% reduction in the PER.

V. Conclusion

In this work, we analysed several error correction techniques for a connectionless wireless system based on BLE. The overall aim was to provide reliable communication and maximise the energy efficiency of the transmitter. Using real BLE statistics collected in office and residential environments, each technique was justified and validated on its own. We then combined all techniques and designed a distributed system, with the emphasis being put on maximising the reliability, minimising the amount of traffic and spreading the computational burden. Using data collected in a typical 2-storey

house, the performance of the designed system was verified. By introducing two orthogonally polarised antennas at the receiver, the PER was reduced by 10%. By having three receivers throughout the house and combining traffic from all of them in a central hub, the packet loss was further reduced by 54%. Error correction using the existing CRC code corrected 20% of the remaining packets. All three techniques do not alter the transmitter and hence do not compromise its energy efficiency. Finally, application-layer coding, even at its simplest form, enabled significant reduction in the overall PER. Although this time the transmitter's energy was spent on encoding packets, it was shown that the resulting energy loss is not significant. In total, based on a two-hour dataset, the overall PER was reduced by 80%.

References

[1] Bluetooth SIG, "Specification of the Bluetooth System - Covered Core Package version: 4.0," 2010.
[2] G. Valadon, F. Le Goff, and C. Berger, "A Practical Characterization of 802.11 Access Points in Paris," in *Proc. 5th Adv. Int. Conf. Telecomm. (AICT)*, 2009, pp. 220–225.
[3] A. Akella, G. Judd, S. Seshan, and P. Steenkiste, "Self-management in Chaotic Wireless Deployments," in *Proc. 11th An. Int. Conf. on Mobile Comp. and Networking (MobiCom)*. ACM, 2005, pp. 185–199.
[4] R. Vaughan, "Polarization diversity in mobile communications," *IEEE Trans. Veh. Technol.*, vol. 39, no. 3, pp. 177–186, 1990.
[5] E. Tsimbalo, X. Fafoutis, and R. Piechocki, "CRC Error Correction for Energy-Constrained Transmission," in *Proc. 26th IEEE Int. Symp. Personal, Indoor and Mobile Radio Commun. (PIMRC)*, 2015.
[6] E. Tsimbalo, X. Fafoutis, and R. Piechocki, "Fix It, Dont Bin It! - CRC Error Correction in Bluetooth Low Energy," in *Proc. 2nd World Forum on Internet of Things (WF-IoT)*, 2015.
[7] V. Kumar and O. Milenkovic, "On graphical representations of algebraic codes suitable for iterative decoding," *IEEE Commun. Lett.*, vol. 9, no. 8, pp. 729–731, 2005.
[8] T. K. Moon, *Error Correction Coding: Mathematical Methods and Algorithms*. Wiley-Interscience, 2005.
[9] J. Feldman, M. Wainwright, and D. Karger, "Using Linear Programming to Decode Binary Linear Codes," *IEEE Trans. Inf. Theory*, vol. 51, no. 3, pp. 954–972, Mar. 2005.
[10] S. Boyd, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2010.
[11] S. Barman, X. Liu, S. Draper, and B. Recht, "Decomposition methods for large scale LP decoding," *2011 49th Annu. Allert. Conf. Commun. Control. Comput.*, vol. 59, no. 12, pp. 253–260, 2011.
[12] P. Woznowski, X. Fafoutis, T. Song, S. Hannuna, M. Camplani, L. Tao, A. Paiement, E. Mellios, M. Haghighi, N. Zhu, G. Hilton, D. Damen, T. Burghardt, M. Mirmehdi, R. Piechocki, D. Kaleshi, and I. Craddock, "A Multi-modal Sensor Infrastructure for Healthcare in a Residential Environment," in *Proc. Int. Conf. Communications Workshops (ICCW)*, Jun. 2015, pp. 271–277.
[13] M. W. Abdullah, X. Fafoutis, E. Mellios, M. Klemm, and G. Hilton, "Investigation into Off-Body Links for Wrist Mounted Antennas in Bluetooth Systems," in *Proc. Loughborough Antennas and Propagation Conf. (LAPC)*, 2015.
[14] E. Mellios, A. Goulianos, S. Dumanli, G. Hilton, R. Piechocki, and I. Craddock, "Off-body Channel Measurements at 2.4 GHz and 868 MHz in an Indoor Environment," in *Proc. 9th Int. Conf. on Body Area Networks (BODYNETS)*, 2014.