

Energy Savings via Harnessing Partial Packets in Body Area Networks

Georgios Angelopoulos
MIT
Cambridge, MA, USA
georgios@mit.edu

Anantha Chandrakasan
MIT
Cambridge, MA, USA
anantha@mit.edu

Muriel Médard
MIT
Cambridge, MA, USA
medard@mit.edu

ABSTRACT

This work considers the incorporation, implications and potential energy savings of partial packet recovery schemes in Body Area Networks (BANs). Received packets which have not been fully corrected by the physical layer, called *partial*, are discarded by the vast majority of BAN protocols, as opposed to *valid* packets, which satisfy the error detection check and are propagated to higher layers. In typical networks using ARQ protocols, dropping partial packets results in retransmissions. However, because these packets contain useful information, partial packet recovery schemes have been proposed with demonstrated throughput and reliability benefits, targeting mostly wireless LANs. In order to quantify the potential energy benefits of harnessing partial packets in BANs, we use an experimental setup with four sensors mounted on a human body, transmitting information to a receiving node in a typical office environment. By precisely modeling the state transitions and energy consumption of sensors, we compare the efficiency of a baseline ARQ protocol against a scheme which leverages information in partial packets. Our results indicate that exploiting partial packets reduces on average the energy consumption of our sensors by 8-20%. The energy savings are pronounced in challenged channel conditions of high PER, where they can be up to 50%.

Keywords

partial packets reception, wireless sensor networks, energy efficient communications, channel coding, erasure coding

1. INTRODUCTION

During the last years, technological advances in microelectronics and novel wireless communications schemes enabled the design and successful deployment of numerous low power wireless sensor networks (WSNs), an emerging research and industrial area of networking. Their wide range of applications include body area networks (BANs), a promising sub-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

BODYNETS 2014, September 29-October 01, London, Great Britain
Copyright © 2014 ICST 978-1-63190-047-1
DOI 10.4108/icst.bodynets.2014.258226



Figure 1: A typical BAN with sensors monitoring vital health signals and transmitting them through intermediate devices to destinations for assessment and storage.

class of WSNs with stringent energy budget and data reliability constraints. Sensor nodes in BANs are typically powered by small energy sources, monitoring vital health signals and transmitting them, usually through an intermediate device, to a collection hub for storage or processing and health assessment, as shown in Fig. 1.

Because of the scarce energy resources, sensor nodes in BANs make use of several techniques to achieve energy efficient data communications. One of them is the use of forward error correction (FEC) schemes at the physical layer (PHY)[17, 13, 15]. Intelligently inserting redundancy by transforming an uncoded message of k bits into a coded packet of n bits increases the probability of the receiver to correctly decode and recover transmitted information. An alternative approach for enhanced energy efficiency and data reliability is the use of erasure codes [18, 19]. Erasure codes can be utilized to introduce cross-packet redundancy, transforming a set of K uncoded packets to a set of N coded packets. Apart from PHY FEC and erasure codes, several cross-layer techniques have been proposed in the literature to overcome shortcomings of single layer approaches [20, 5].

Although all the aforementioned techniques significantly improve on average the data reliability, under certain conditions, they fail to recover the initial information. In order to prevent erroneous packets to be forwarded to higher layers of the protocol, error detection codes (e.g. cyclic redundancy checks-CRCs) are typically inserted between the recovery schemes and the application layer, completely discarding packets with erroneous bits. In the rest of the paper, packets which satisfy the error detection rule and can

be propagated to higher layers are called *valid*, while packets which have been processed by the PHY and contain some erroneous bits are called *partial*.

In principle, exact knowledge of the channel state information combined with continuous power, rate and modulation adaptation could guarantee minimum overhead and error free transmission. However, the highly varying nature of the communications medium in BANs, due to intrinsic noise, external interfering signals and body movements, results in significant *SNR* variations and makes quite challenging the task of channel tracking [6]. In addition, typical BANs sensor nodes can not afford energy or computational intensive continuous channel sensing and probing mechanisms. Thus, when the channel quality drops, sensors' PHY might fail to correct all errors in received packets, resulting in partial packets. Majority of BANs using ARQ protocols discard these packets and request a retransmission, which in some cases results in increased total energy consumption.

In order to improve the "all-or-nothing" nature of the conventional BANs protocols and reduce the number of retransmissions, partial packet recovery methods can be used. These methods harness packets which have not fully corrected by the PHY, increasing network throughput and resource utilization, targeting mostly WLANs [8, 16, 3, 14, 21]. This work examines the applicability of partial packet recovery mechanisms in BANs and the potential energy savings, extending sensor nodes' battery lifetime. Experiments are carried out with four wireless sensor nodes mounted on a human body in a typical office environment. A specific partial packet recovery method suitable for BANs sensors is briefly presented and its incorporation in a realistic protocol is discussed. By precisely modeling the power consumption of sensors, we demonstrate the energy savings achieved by harnessing partial packets. In challenged channel conditions of high PER, exploitation of information contained in partial partials can significantly reduce the number of retransmitted packets, resulting in up to 50% energy reduction, while a 8-20% energy savings is observed on average across all channel conditions.

2. HARNESSING PARTIAL PACKETS

2.1 Background and Prior Work

The potential of exploiting correct information within partial packets has attracted significant interest in the past and several mechanisms have been presented in the literature. Authors in [8] were among the first who identified the benefits of extracting the correct information and combining partial packets to reduce the total number of transmitted packets in multi-hop sensor networks, leveraging the opportunistic nature of the wireless medium. A similar scheme was presented in [9], increasing the channel utilization in sensor networks by concatenating blocks and allowing larger packets to be transmitted. Each block, having its own CRC, could be separately checked at the receiver for its validity.

Although these schemes emphasized the benefits of combining partial packets in sensor networks, they are associated with some overhead which, under certain conditions, may result in less energy efficient solutions. For instance, inserting multiple CRCs per packet or indicating within the feedback information which parts of the packets were corrupted, increases the total transmitted and received data, respectfully.

An alternative way to distinguish correct and erroneous information within a partial packet is to make use of soft PHY information. In this approach, every decoded bit by the PHY is annotated with a reliability metric, enabling the link or higher layers to identify which parts of the packet are more likely to contain erroneous bits [14, 21]. HARQ schemes operate in a similar manner, accumulating soft PHY information across multiple retransmissions [7]. Unfortunately, these schemes can only be implemented on software defined radio platforms and not over commodity radios, since they violate the conventional abstraction principle between the layers of the protocol stack. PHY independent schemes have also been proposed [16, 10], mainly targeting increased throughput in wireless LANs. These schemes take advantage of error detection and correction capabilities of coding techniques, such as Reed-Solomon codes. By encoding the transmitted data and imposing an algebraic structure, the receiver can detect which parts of a partial packet contain no errors, leveraging these packet chunks.

2.2 BANs and Partial Packets

Considering the unique constraints and characteristics in BANs, a partial packet recovery scheme should ideally have the following features:

- introduce zero fixed transmission overhead, so that when packets are correctly received, no extra information is unnecessarily associated with them,
- not increase the feedback information, because reception energy is a considerable component of the total energy consumption,
- be PHY independent and easily applicable to existing platforms and devices, and
- introduce minimal additional complexity on the sensor side, pushing the complexity on the more computationally capable receivers.

Among the partial packet recovery schemes presented in the literature, Packetized Rateless Algebraic Consistency (PRAC) scheme satisfies all aforementioned requirements and can be tailored for easy incorporation in BANs [3]. PRAC is a PHY independent recovery scheme, which can be implemented as a software patch in an existing network stack. It does not introduce any fixed transmission overhead, such as multiple CRCs or pilot bits within a packet, and its rateless encoding process requires no extra feedback information and minimal additional complexity.

2.2.1 Encoding Process of PRAC

Assume that P_1, P_2, \dots, P_k are packets to be transmitted, each of them with L bits. By segmenting packet P_i into l blocks (p_{ij} , where $j = 1, \dots, l$), each block contains $L/l = q$ bits and can be considered as an element from a finite field $\mathbb{F}(2^q)$. The encoding process is equivalent of multiplying the initial packets with a randomly generated matrix, introducing a cross-packet algebraic structure, similarly to random linear network coding (RLNC) [11]. This structure is exploited at the receiver to enable the identification of correct and erroneous information. A block of an encoded packet P'_i is calculated as:

$$p'_{ij} = \sum_{i=1}^l c_i p_{ij}, \quad (1)$$

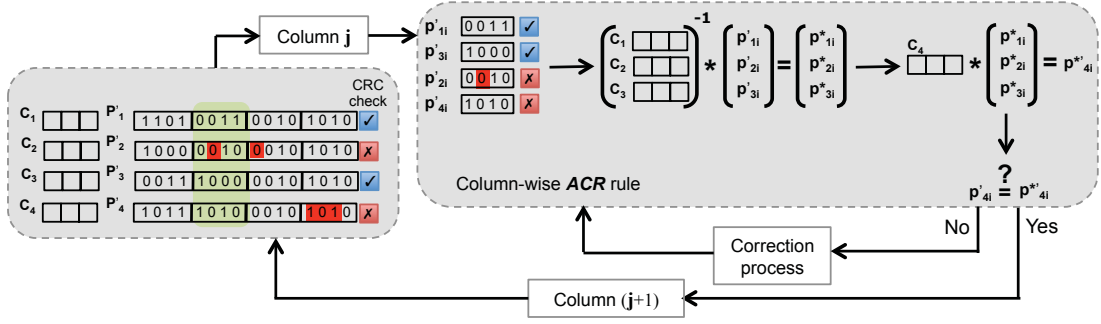


Figure 2: PRAC’s iterative decoding process operates sequentially over columns of the buffered packets’ matrix (P_1' to P_4'), using the ACR check rule to examine the consistency of a column and the correction process to recover erroneous blocks.

where $c_i \in \mathbb{F}(2^q)$ is a randomly selected coefficient. It should be noted here that encoded packets are of the same length as the uncoded ones, and that the coefficient values can be generated in both the transmitter and receiver from a pseudo-random number generator, seeded with the same initial value. This ensures that no extra overhead is transmitted. The energy analysis and requirements of the encoding process implemented in a low power accelerator is presented in [4], implying that the required energy to perform these operations is negligible because the energy per bit for the encoding process is significantly lower (approximately two orders of magnitude) than the energy per bit for transmission or reception.

2.2.2 Decoding Process of PRAC

PRAC buffers valid and partial packets and initiates its decoding process when the number of buffered packets exceeds k . A sequential, column-wise process attempts to identify correct packet chunks, and when erroneous parts are located, a correction process attempts to recover their initial values. Algebraic consistency rule (ACR) is a check which examines if the inserted algebraic structure exists over the currently considered packets’ column. If the result is positive, then with very high probability the initial packet blocks can be recovered and the decoding process proceeds to the next column. If the result of the ACR check is negative, then inconsistency exists in the considered column because of erroneous blocks. An iterative correction process is triggered in that case, consisted of an optimized search algorithm and multiple ACR checks. An example of the PRAC’s decoding process is shown in Fig. 2, in which $k = 3$, $L = 16$ and $q = 4$.

3. EXPERIMENTAL MEASUREMENTS OF PARTIAL PACKETS IN BANS

Because of the strict energy constraints and the regulatory requirements, sensors in BANS transmit their signals in low output power levels, several orders of magnitude lower than typical wireless LANs and cellular networks. In addition to the external interference, the highly dynamic nature of wireless medium, mainly due to movements and specific positions of the human body, causes the effective channel SNR to significantly vary over time [12]. Channel variations challenges the energy efficiency of data transmission in BANS. This often results in corrupted packets, which can

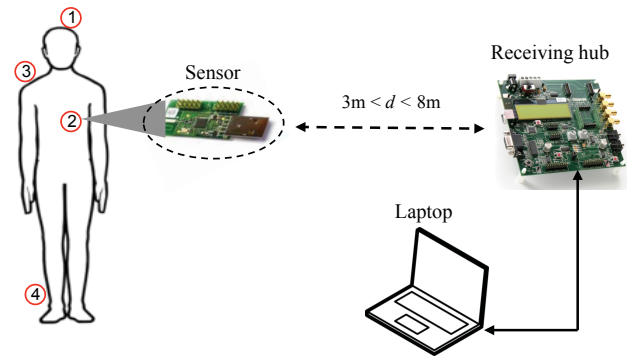


Figure 3: Experimental setup: four on-body sensors transmit information to a receiving hub. Collected data traces are stored in a laptop for further processing.

not be fully recovered by the PHY FEC correction process, resulting in requests for retransmissions. Thus, exploiting partial packets and reducing the number of retransmitted packets could be a major enabling factor in improving the energy efficiency of BANS.

3.1 Experimental Setup

In order to quantify the frequency of occurrence and energy overhead associated with partial packets in a typical BAN application, we make use of the experimental setup shown in Fig. 3. Four on-body transmitting sensors are communicating with a receiving hub, mounted on a wall and connected to a PC storing the collected data traces and processing them offline. The on-body sensors are equipped with CC2511 transceivers [2] transmitting an ECG waveform stored in the on-board memory, and the receiving hub is a CC2500 development board [1]. The sensors are placed in four different body locations: *i*) top of scalp, *ii*) left chest, *iii*) back, and *iv*) right ankle. Scheduling among the sensors is performed by the hub, allocating time slots in each sensor, ensuring that no cross-sensors interference exists.

The output transmission power (P_{TXout}) of the sensors is adjusted from $-25dBm$ to $1dBm$. The transmission rate is fixed at $250kbps$ and communication is performed at the $2.4GHz$ ISM frequency band. 2-FSK modulation and coher-



Figure 4: Packet format used in our experiments.

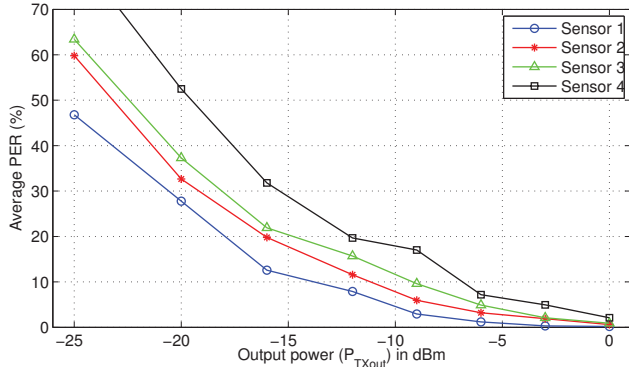


Figure 5: Average packet error rate with respect to the transmission output power.

ent demodulation, a short interleaver, a convolutional rate 1/2 code and a hard Viterbi decoder are used at the PHY. The packet format is shown in Fig. 4. Upon the reception of a packet by the hub, its CRC status is examined and it is buffered on the PC as valid or partial, depending on the result of the CRC check. The packet encoding described by Eq. 1 is performed by the microcontroller of the sensor, while the decoding process of the partial packet recovery mechanism is implemented in software on the PC.

All experiments and measurements are performed in a typical indoors campus environment, with no control over the nearby interfering networks, e.g. WiFi. The distance between the sensors and the receiving hub varies in the 3–8m range as the person with the mounted sensors is moving in the room, performing typical tasks, including sitting in a chair.

3.2 Data Measurements

Extensive channel modeling and experimental measurements have been published in the literature [6], capturing the detailed characteristics of the wireless medium around the human body in indoors and outdoors environments. In this work, we provide some experimental measurements of typical low power sensors communicating around the human body, not as an exhaustive modeling effort, but more as a characterization of the channel quality experienced during our measurements.

Fig. 5 plots the average packet error rate (PER) for each P_{TXout} value in the four links of our experimental setup. As expected, the PER increases as P_{TXout} decreases. All four links exhibit approximately similar channel characteristics, with sensor 1 achieving the best average performance because of its position and the almost always available line-of-sight path to the receiver. Fig. 6 plots the variations on the channel quality due to channel impairments and body movements for one link of our experimental setup. Error bars for all links are not included for readability purposes.

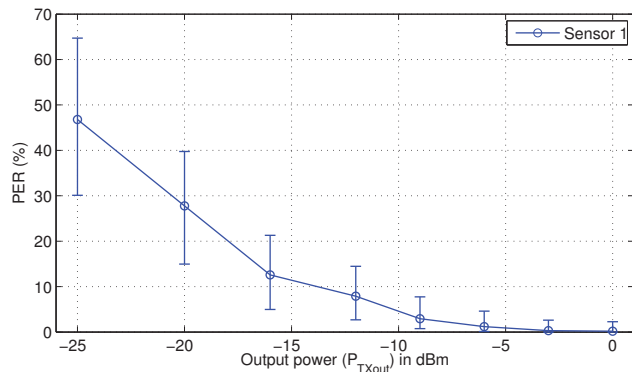


Figure 6: Channel quality variation of the link ‘Sensor1-hub’; similar behavior is observed in the rest links but not included for readability purposes.

Table 1: Power consumption of a sensor in different states.

State	Symbol	Value
Transmission	P_{TX}	78mW, $P_{TXout} = 0dBm$
		54.5mW, $P_{TXout} = -6dBm$
		48mW, $P_{TXout} = -12dBm$
		42mW, $P_{TXout} = -20dBm$
Reception	P_{RX}	63mW
Idling	P_{IDLE}	14.4mW
Sleep	P_{SLEEP}	$\sim 0W$

4. PERFORMANCE EVALUATION

Using the collected data traces from the experimental setup, the energy benefits associated with harnessing partial packets in BANs are explored. We consider the energy comparison of two scenarios: one using a baseline ARQ protocol discarding partial packets and another using PRAC.

4.1 Energy Consumption Modeling

In our analysis, we consider the energy consumption of the sensors only, without taking into account the energy of the receiving hub. We assume the receiving node is a more powerful and less constrained device compared to BANs sensors. This is a typical assumption in majority of asymmetric networks. For the energy consumption calculations, we assume that sensors are transitioning among four different states: transmission, reception, idling and sleep states. The power consumption associated with each state is shown in Table 1 [2]. As expected, the power consumption during the transmission state depends on P_{TXout} . In the rest of our analysis, we assume that the power consumption during sleep states is approximately zero and sensors do not consume any energy.

The upper part of Fig. 7 shows the sequence of transmitted and received packets by a sensor, with timing details, which applies in both scenarios of using the baseline ARQ and PRAC scheme. Transmission of a packet is followed, after an interframe space interval (t_{sifs}), by an acknowledgment packet (ACK) sent by the hub, indicating its successful reception or not. The transition timings between the different states (t_1, t_2, t_3, t_4 and t_{sifs}) and the packet (ACK) transmission (reception) durations are summarized in Table 2 [1, 2]. In the lower part of Fig. 7 the power consumption

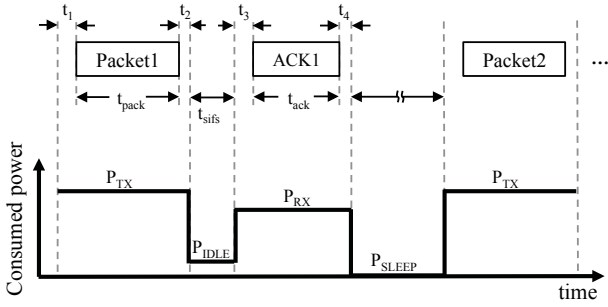


Figure 7: State transition and power diagram for a sensor.

Table 2: Timing notation and their values.

Symbol	Description	Value
t_1	transition time from SLEEP to TX mode	0.9ms
t_2	transition time from TX to IDLE mode	1 μ s
t_3	transition time from IDLE to RX mode	0.3ms
t_4	transition time from RX to SLEEP mode	1 μ s
t_{sifs}	interframe time interval	0.4ms
t_{frame}	transmission time of a packet	4.45ms
t_{ack}	reception time of an ack	0.7ms

of a sensor is shown over time. We assume that during the transition times, i.e. sleep to transmission mode, reception to idle mode, etc., the sensor has a constant power consumption.

4.2 Energy Savings from Harnessing Partial Packets

According to our energy model, the expected energy consumption of a sensor (\bar{E}_{sen}), given a transmission output power, in order to transmit N data packets to the hub is:

$$\bar{E}_{sen}(N) \Big|_{P_{TXout}} = \mathbb{E}[NE_{pack}] \Big|_{P_{TXout}} = \bar{N}E_{pack} \Big|_{P_{TXout}}, \quad (2)$$

where \bar{N} is the expected number of transmitted packets and received ACKs, including the retransmitted ones, and E_{pack} is the energy of the sensor to transmit a packet at P_{TXout} and receive its ACK. \bar{N} depends on the transmission parameters, e.g. P_{TXout} , experienced channel, receiver's sensitivity and processing capabilities, such as use of partial packet recovery methods or not. According to Fig. 7, E_{pack} is:

$$E_{pack} = P_{TX}(t_1 + t_{pack} + t_2) + P_{IDLE}t_{sifs} + P_{RX}(t_3 + t_{ack} + t_4). \quad (3)$$

Examining Eq. 2 and 3, \bar{E}_{sen} depends linearly on P_{TX} ; however, decreasing P_{TX} impairs packets' received signal quality, increasing the PER and consequently \bar{N} . Thus, a trade-off exists and careful optimization is required to ensure minimum energy consumption. Fig. 8 plots the total energy consumption of the four sensors as they transmit N packets to the hub and receive the corresponding ACKs, using ARQ and PRAC schemes, respectively. In this figure,

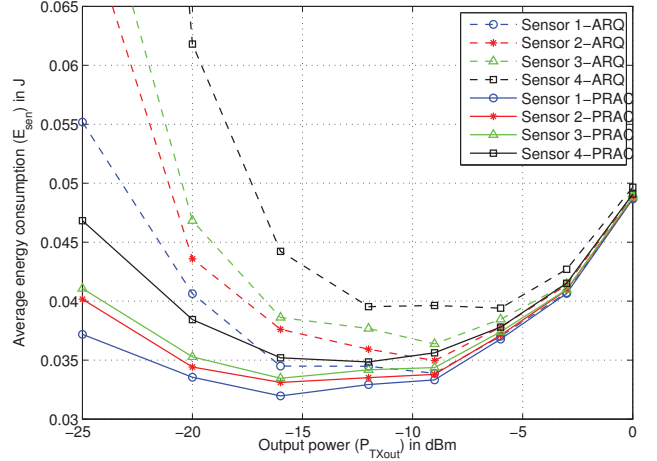


Figure 8: Total energy consumption of sensors using PRAC, harnessing partial packets, for transmission of $N = 100$ packets to the hub, with regards to P_{TXout} .

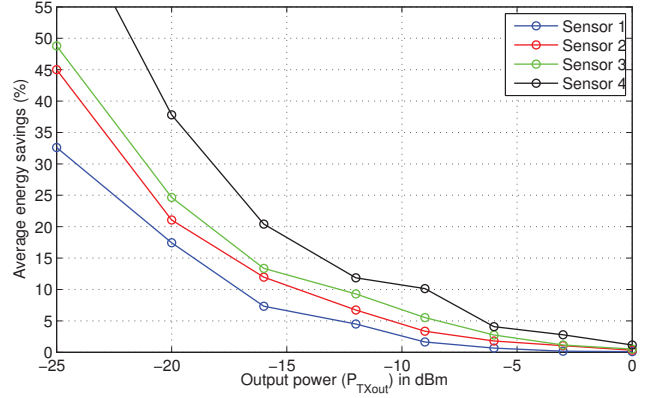


Figure 9: Average energy savings by harnessing partial packets.

we assume $N = 100$. When P_{TXout} is high enough, PER is low and \bar{E}_{sen} scales with it. However, decreasing P_{TXout} below some value significantly impairs quality of transmission, increasing PER and resulting in excessive retransmissions. The energy overhead of these retransmitted packets exceeds the energy savings through scaling P_{TXout} , increasing the expected energy consumption and giving the 'U-shape' in all lines of Fig. 8.

Harnessing partial packets with PRAC decreases the required number of total transmitted packets and results in lower energy consumption. In Fig. 9, the average energy savings are plotted for each sensor with respect to P_{TXout} . As expected, the benefits are pronounced in the moderate to high PER regime (medium to low P_{TXout}) and can be up to 50% in challenged channel conditions. Averaging the energy savings of each sensor in our experimental setup across all values of P_{TXout} range, a 8-20% energy reduction is observed with PRAC.

5. CONCLUSIONS

Body Area Networks (BANs) has been an emerging research field with the potential to revolutionize several applications, e.g. health care, sports training, etc. Data communications is usually a significant component of the total energy consumption in a typical BAN sensor thus, optimizing data transmission and reception is of utmost interest. Although use of PHY FEC and cross-layer schemes improves data reliability, the “all-or-nothing” nature of prevailing BANs protocol stacks, similarly to the vast majority of modern wireless systems using ARQ protocols, discards partial packets and retransmissions are requested, usually resulting in excessive energy consumption.

In this work, an experimental setup, with four commercial sensors mounted on a human body and a receiving node, is used to examine and quantify the inefficiency caused by dropping partial packets. All experiments are performed indoors, in a typical office environment, and collected data traces are stored in a PC for offline processing. Among the schemes proposed in the literature, PRAC is selected as the most appropriate partial packet recovery mechanism, because it does not introduce any fixed transmission overhead, it does not modify or increase the feedback mechanism, it has minimal encoding complexity and, most importantly, it is PHY independent and can be easily inserted in existing wireless systems as a software patch.

After a detailed energy modeling of the transmission and reception process in a sensor, our results reveal that exploiting partial packets have the potential to offer significant energy savings. In a challenged link with high PER, the energy benefits can be up to 50% compared to a baseline ARQ scheme, while a 8-20% energy reduction is observed for each sensor on average across all P_{TXout} values.

6. ACKNOWLEDGMENTS

This material is based upon work supported by VT iDirect, Inc. under award No. 020903-001.

7. REFERENCES

- [1] Texas instruments, CC2500dk Data Sheet. <http://www.ti.com/tool/cc2500-cc2550dk>.
- [2] Texas instruments, CC2511 Data Sheet. <http://www.ti.com/lit/gpn/cc2511f32>.
- [3] G. Angelopoulos, A. Chandrakasan, and M. Médard. PRAC: Exploiting partial packets without cross-layer or feedback information. In *Communications (ICC), 2014 IEEE International Conference on*, June 2014.
- [4] G. Angelopoulos, M. Médard, and A. P. Chandrakasan. Energy-aware hardware implementation of network coding. In *Proceedings of the IFIP TC 6th International Conference on Networking, NETWORKING'11*, pages 137–144, 2011.
- [5] G. Angelopoulos, A. Paidimarri, A. Chandrakasan, and M. Médard. Experimental study of the interplay of channel and network coding in low power sensor applications. In *Communications (ICC), 2013 IEEE International Conference on*, pages 5126–5130, June 2013.
- [6] R. Cavallari, F. Martelli, R. Rosini, C. Buratti, and R. Verdone. A survey on wireless body area networks: Technologies and design challenges. *Communications Surveys Tutorials, IEEE*, PP(99):1–23, 2014.
- [7] D. Chase. Code combining—a maximum-likelihood decoding approach for combining an arbitrary number of noisy packets. *Communications, IEEE Transactions on*, 33(5):385–393, May 1985.
- [8] H. Dubois-Ferrière, D. Estrin, and M. Vetterli. Packet combining in sensor networks. In *SenSys*, 2005.
- [9] R. K. Ganti, P. Jayachandran, H. Luo, and T. F. Abdelzaher. Datalink streaming in wireless sensor networks. In *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems, SenSys '06*, pages 209–222, 2006.
- [10] B. Han, A. Schulman, F. Gringoli, N. Spring, B. Bhattacharjee, L. Nava, L. Ji, S. Lee, and R. Miller. Maranello: Practical partial packet recovery for 802.11. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation, NSDI'10*, pages 14–14, 2010.
- [11] T. Ho, M. Medard, R. Koetter, D. Karger, M. Effros, J. Shi, and B. Leong. A random linear network coding approach to multicast. *Information Theory, IEEE Transactions on*, 52(10):4413–4430, Oct 2006.
- [12] H. Hongwei, S. Wei, X. Youzhi, and Z. Hongke. The effect of human activities on 2.4 ghz radio propagation at home environment. In *Broadband Network Multimedia Technology, 2009. IC-BNMT '09. 2nd IEEE International Conference on*, pages 95–99, Oct 2009.
- [13] S. L. Howard, C. Schlegel, and K. Iniewski. Error control coding in low-power wireless sensor networks: when is ECC energy-efficient? *EURASIP J. Wirel. Commun. Netw.*, 2006(2):29–29, Apr. 2006.
- [14] K. Jamieson and H. Balakrishnan. PPR: partial packet recovery for wireless networks. In *SIGCOMM*, Aug. 2007.
- [15] Z. Kashani and M. Shiva. Power optimised channel coding in wireless sensor networks using low-density parity-check codes. *Communications, IET*, 1(6):1256–1262, Dec. 2007.
- [16] K. C.-J. Lin, N. Kushman, and D. Katabi. Ziptx: Harnessing partial packets in 802.11 networks. In *Proceedings of the 14th ACM International Conference on Mobile Computing and Networking, MobiCom '08*, pages 351–362, 2008.
- [17] S. Lin and D. J. Costello. *Error Control Coding: Fundamentals and Applications*. Prentice Hall: Englewood Cliffs, NJ, second edition, 2004.
- [18] M. Luby, M. Mitzenmacher, M. Shokrollahi, and D. Spielman. Efficient erasure correcting codes. *Information Theory, IEEE Transactions on*, 47(2):569–584, Feb. 2001.
- [19] I. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, 1960.
- [20] M. Vuran and I. Akyildiz. Error control in wireless sensor networks: A cross layer analysis. *Networking, IEEE/ACM Transactions on*, 17(4):1186–1199, Aug. 2009.
- [21] G. R. Woo, P. Kheradpour, D. Shen, and D. Katabi. Beyond the bits: cooperative packet recovery using physical layer information. In *MobiCom*, 2007.