

Iterative Decoding for Redistributing Energy Consumption in Wireless Sensor Networks

R. G. Maunder, A. S. Weddell, G. V. Merrett, B. M. Al-Hashimi and L. Hanzo
ECS, University of Southampton, SO17 1BJ, UK.
Corresponding author: rm@ecs.soton.ac.uk

Abstract—In this paper, we propose a method for desirably redistributing a wireless sensor network’s energy consumption from its sensor nodes (which may have scarce energy resources obtained through energy harvesting, for example) to its central node (which often has an abundant energy resource, such as the mains). At the cost of increasing the central node’s decoding complexity, our method facilitates (1) a significant reduction in the number of times the sensor nodes are required to retransmit data owing to transmission errors and/or (2) a reduction of up to 3.99 dB in the sensor node’s total transmit energy consumption. We show that our approach can reduce the overall energy consumption of transmitting sensor nodes by more than 20% in practice.

I. INTRODUCTION

Since its introduction, the IEEE 802.15.4 standard designed for low-rate Wireless Personal Area Networks (WPANs) [1] has found application in pervasive Wireless Sensor Networks (WSNs). These typically comprise a number of smart sensor nodes that are required to maintain sporadic but reliable communications with each other for extended periods of time. A star-structured network topology is often employed, with all data frames being routed via a central node, which coordinates the reactions of the system’s application layer to the sensed data. Owing to its integration into the higher-level system, the central node typically has abundant energy resources, such as the mains. By contrast, some or all of the sensor nodes may have limited energy resources, relying on energy harvesting for example, to maintain their operation. This unequal distribution of energy in pervasive WSNs motivates this paper’s introduction of an iterative-decoding aided augmentation to the IEEE 802.15.4 2.450 MHz PHYsical layer (PHY) [1], as detailed in Section II.

In Section III of this paper, we show that our augmented PHY facilitates a significantly improved data Frame Error Rate (FER) performance and/or requires a significantly reduced transmit energy. While this is achieved at the cost of an increased decoding complexity (as well as a slightly higher encoding complexity), opting for this trade-off is desirable in the pervasive WSNs described above. More specifically, when transmitting, the limited energy resources of the sensor nodes benefit from the reduced transmit energy consumption that is facilitated by the augmented PHY. Additionally, owing to the reduced FER that this yields, the data frames will become unacknowledged less frequently and less sensor node energy will be consumed by retransmissions. In fact, if we assume that all data frames will be received without error,

further sensor node energy can be saved by not listening for acknowledgements at all, although data can be lost using this approach. Upon the reception of the data frames, the central node can afford the increased decoding complexity owing to its abundant energy resources. When the central node relays the data frames, the standard PHY may be employed without augmentation so that the sensor nodes do not incur the augmentation’s increased decoding complexity. In this case, a low FER can be maintained by employing an increased transmit energy, which can be afforded by the central node owing to its abundant energy resources.

Finally, we offer our conclusions in Section IV.

II. AUGMENTED PHY

In this section, we introduce our augmented PHY. This may be invoked to convey the PHY payloads of the IEEE 802.15.4 data frames [1, Section 5.5.3.2], which have 118 possible transmission frame lengths in the range $M_p \in [10 \dots 127]$ bytes [1, Figure 11]. Note that the proposed augmentation cannot be invoked to convey the $M_h = 6$ -byte headers of the data frames [1, Section 6.3]. This is because the headers contain the synchronisation sequence and the PHY payload length, both of which must be recovered before iterative decoding can commence in the proposed augmentation to the PHY. However, using the standard PHY without augmentation to convey the headers has the benefit of maintaining compatibility with existing IEEE 802.15.4 networks.

Let us begin by summarizing the operation of the standard PHY, when conveying the PHY payloads of data frames, before describing how this is augmented by the proposals of this paper. The PHY employs Direct Sequence Spread Spectrum (DSSS) operation [1, Section 6.5], invoking Pseudo-random Noise (PN) spreading [2] and Offset Quadrature Phase-Shift Keying (O-QPSK) [3], as shown in Figure 1. During PN spreading, the M_p -byte PHY payload \mathbf{b} is decomposed into sets of $k = 4$ consecutive bits, which are mapped to $n = 32$ -chip codewords [1, Table 24]. These codewords are concatenated to obtain the chip sequence \mathbf{c} , which has 118 possible lengths $N_p = 8M_p n/k \in \{640, 704, 768, \dots, 8128\}$, like the PHY payload \mathbf{b} . Finally, O-QPSK modulation is employed to convert the chip sequence \mathbf{c} into the modulated signal \mathbf{s} [1, Section 6.5.2.4], as shown in Figure 1.

During the demodulation of the received signal $\tilde{\mathbf{s}}$, the demodulator expresses its confidence that a particular transmitted chip in \mathbf{c} had a ‘0’- or a ‘1’-value using a corresponding

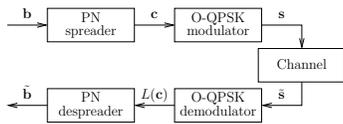


Fig. 1. PHY schematic.

Logarithmic Likelihood Ratio (LLR) [4] in the sequence $L(\mathbf{c})$. Finally, the PN despreader of Figure 1 decodes the LLRs $L(\mathbf{c})$ [5] in order to obtain a reconstruction $\hat{\mathbf{b}}$ of the PHY payload \mathbf{b} .

A schematic for our proposed augmentation to the PHY is provided in Figure 2. Note that this schematic retains the PN spreading and O-QPSK modulation of the standard PHY depicted in Figure 1. However, in the transmitter of Figure 2, the PN spreader is serially concatenated [6] with a rate-1 encoder [7], which is invoked before O-QPSK modulation. As is common for concatenated codes [6], the PN spreader and the rate-1 encoder are separated by an interleaver. Note that the additional processing required to perform interleaving and rate-1 encoding is responsible for the marginally increased encoding complexity of the augmented PHY. The operation and complexity of these components is discussed in Sections II-A and II-B.

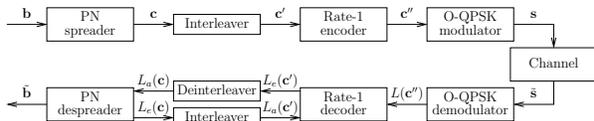


Fig. 2. Augmented PHY schematic.

Iterative decoding [8] is employed in the receiver of the augmented PHY, which repeatedly alternates the operation of the PN despreader and the rate-1 decoder, as shown in Figure 2. This is in contrast to the standard PHY's receiver, which requires only the 'one-shot' non-iterative activation of the PN despreader. Since the augmented PHY invests more decoding complexity than the standard PHY, it can be expected to achieve a lower Payload Error Rate (PER) for a particular transmit energy. The iterative decoding process is described in Section II-C, while the resultant PER improvement, transmit energy gain and decoding complexity are quantified in Section III.

Note that our iterative decoding approach differs from the iteratively decoded DSSS schemes found in the literature [9], [10], which do not invoke the PN despreader during the iterative decoding process. Instead, the PN despreader is either replaced with a low-rate iteratively decoded turbo code [9] or iterative turbo decoding is invoked only after the independent PN despreading operation [10]. Our approach has a greater resemblance to chip-interleaved Code Division Multiple Access (CDMA) schemes [11], which employ iterative despreading and MultiUser Detection (MUD) [12].

In the following sub-sections, we detail the interleaving, rate-1 encoding and iterative decoding processes employed by

the augmented PHY of Figure 2.

A. Interleaving

In the transmitter of the augmented PHY of Figure 2, the N_p -chip sequence $\mathbf{c} = \{c[n]\}_{n=1}^{N_p}$ is interleaved, rather than O-QPSK modulated like in the standard PHY. Here, the interleaver re-arranges the order of the chips in the sequence \mathbf{c} in the particular manner described by the interleaver mappings $\mathbf{\Pi} = \{\Pi[n]\}_{n=1}^{N_p}$. More specifically, the resultant interleaved chip sequence $\mathbf{c}' = \{c'[n]\}_{n=1}^{N_p}$ is obtained according to the mappings $c'[n] = c[\Pi[n]]$, where $\Pi[n] \in [1, N_p] \neq \Pi[m] \in [1, N_p] \forall n \in [1, N_p] \neq m \in [1, N_p]$. Similarly, the deinterleaver in the augmented PHY's receiver of Figure 2 is employed to reverse this re-arrangement and hence restore the original ordering.

The above-mentioned re-arrangement operations are necessary because the achievable iterative decoding performance and, hence, the achievable PER performance are commensurate with the degree to which the chip ordering in \mathbf{c}' appears to be randomized [13]. Note that an improved PER performance can be expected for longer PHY payloads, since the interleaver's ability to pseudo-randomly re-arrange the order of the chips in the sequence \mathbf{c}' is commensurate with its length N_p . As this implies, a different set of interleaver mappings $\mathbf{\Pi}$ is required for each of the 118 possible lengths $N_p \in \{640, 704, 768, \dots, 8128\}$ of the chip-sequence \mathbf{c} . These mappings must be available in both the transmitter and the receiver of the augmented PHY, requiring $\sum_{N_p} N_p \lceil \log_2(N_p) \rceil \approx 800$ KB of Read Only Memory (ROM). While this can be afforded in the WSN's central node, it is undesirable (and it turns out unnecessary) in a simple sensor node.

This motivates the use of so-called deterministic interleavers [13] in the augmented PHY. Rather than using raw interleaver mappings $\mathbf{\Pi}$, deterministic interleavers require the storage of only a small number of parameters, which unambiguously describe a particular way of re-arranging the chips in the sequence \mathbf{c} . Deterministic interleavers therefore require less ROM than random interleavers. However, it becomes a challenge to find parameters that yield beneficial pseudo-random re-arrangements.

We elected to employ Dithered Relative Prime (DRP) interleaver designs [14], since these are capable of yielding pseudo-random re-arrangements. The parameters of the interleavers were selected using a specially designed Genetic Algorithm (GA) [15], which searched for designs maximising a 'randomness' metric. We found that the resultant designs may be stored using about 12 KB of ROM, which is significantly less than the 800 KB required by the randomly designed interleavers. In Section III, we shall show that there is no PER performance penalty associated with employing our deterministic interleaver designs instead of the randomly designed interleavers.

A further benefit of using DRP interleavers is that they impose only a small additional complexity upon the transmitter of the augmented PHY, since they employ only three low-complexity re-arrangement operations [14]. We may assume

that each of these operations can be implemented using a single clock cycle for each of the N_p chips in the data payload, requiring a total of $3N_p$ clock cycles. This is a conservative estimate, since a practical implementation could process a number of chips in parallel during each clock cycle. The associated energy consumption associated with this processing shall be investigated in Section III-B.

B. Rate-1 encoding

Following chip interleaving in the augmented PHY's transmitter of Figure 2, the resultant N_p -chip sequence \mathbf{c}' is encoded using the rate-1 encoder [7] in order to obtain the sequence $\mathbf{c}'' = \{c''[n]\}_{n=1}^{N_p}$. The rate-1 encoder operates on the basis of a single modulo-2 memory element and a single modulo-2 addition according to $c''[n] = c'[n] \oplus c''[n-1]$, where $n \in [1 \dots N_p]$ and $c''[0] = 0$.

Owing to the simplicity of rate-1 encoding, it imposes only a small additional complexity upon the transmitter of the augmented PHY. In Section III-B, we shall conservatively assume that a single clock cycle is required to encode each of the N_p chips in the data payload. With reference to Section II-A, a total of $4N_p$ clock cycles is therefore required to perform interleaving and rate-1 encoding.

Following rate-1 encoding in the augmented PHY of Figure 2, the encoded chip sequence \mathbf{c}'' is O-QPSK modulated, similarly to the standard PHY of Figure 1. Note that since the encoded chip sequence \mathbf{c}'' contains the same N_p number of chips as the sequence \mathbf{c} , the PHY employing the proposed augmentation transmits the same amount of information as the standard PHY of Figure 1. The direct comparison of these PHYs is therefore fair and we shall employ the PHY without augmentation as a bench-marker in order to assess the performance of the PHY employing the proposed augmentation.

C. Iterative decoding

In the receiver of the augmented PHY shown in Figure 2, the O-QPSK demodulator [4] generates the LLR sequence $L(\mathbf{c}'')$ to express its confidence in the values of the transmitted chips in \mathbf{c}'' . Following this, iterative decoding proceeds with the alternated iterative activation of the rate-1 decoder and the PN despreader. These exchange any new information that they can obtain in the form of the so-called extrinsic LLR sequences $L_e(\mathbf{c}')$ and $L_e(\mathbf{c})$ [16]. After interleaving or deinterleaving as appropriate, these extrinsic LLR sequences are employed as so-called *a priori* LLR sequences $L_a(\mathbf{c})$ and $L_a(\mathbf{c}')$ [16] in order to assist the operation of the other constituent decoder.

Note that at the commencement of iterative decoding, the PN despreader will not yet have been invoked and hence the *a priori* LLR sequence $L_a(\mathbf{c}')$ will be unavailable to assist the operation of the rate-1 decoder. In this case, the extrinsic LLR sequence $L_e(\mathbf{c}')$ is generated by considering only the LLR sequence $L(\mathbf{c}'')$ provided by the O-QPSK demodulator of Figure 2. However, in all subsequent decoding iterations, the rate-1 decoder generates the extrinsic LLR sequence $L_e(\mathbf{c}')$ by considering both the *a priori* LLR sequence $L_a(\mathbf{c}')$ and the LLR sequence $L(\mathbf{c}'')$ provided by the O-QPSK demodulator.

Upon each successive decoding iteration, the rate-1 decoder and the PN despreader obtain more and more confidence in the values of the transmitted chips in \mathbf{c} and \mathbf{c}' . This extrinsic information exchange continues until no more new information about the transmitted chips can be gleaned, whereupon convergence is deemed to have been achieved. This event may be detected by employing the averaging method of [17] to measure the mutual information [18] of the extrinsic LLR sequence $L_e(\mathbf{c}')$ and by comparing it with that measured in the previous decoding iteration. When the mutual information of the extrinsic LLR sequence $L_e(\mathbf{c}')$ stops increasing, iterative decoding can be curtailed and the reconstructed PHY payload $\tilde{\mathbf{b}}$ may be output.

In the receiver of the proposed augmentation to the PHY shown in Figure 2, the PN despreader operates on the basis of the Soft-In Soft-Out (SISO) decoder of [4], while the rate-1 decoder applies the Bahl-Cocke-Jelinek-Raviv (BCJR) algorithm [19] to a trellis structure [20]. In both cases, all calculations are performed in the logarithmic domain, using an eight-entry lookup table to correct the Jacobian approximation [16]. As a result, all calculations can be performed using only Add, Compare and Select (ACS) operations, which may be performed in a single clock cycle by a simple fixed-point Arithmetic and Logic Unit (ALU).

III. RESULTS

In this section we compare and discuss the PERs that may be obtained when using both the augmented PHY and the standard PHY without augmentation to convey PHY payloads comprising various numbers of chips in the range of $N_p = 8M_p n/k \in \{640, 704, 768, \dots, 8128\}$. We also investigate the error rate associated with the PHY headers, which always comprise a total of $N_h = 8M_h n/k = 384$ chips [1, Section 6.3]. As described in Section II, the PHY headers are always transmitted using the PHY without augmentation, even when the augmentation is employed during the transmission of the PHY payload. For this reason, we only consider the Header Error Rate (HER) associated with the standard PHY. Later, the HER and PERs are combined to obtain an expression for the FER. We also quantify the additional encoding and decoding complexity that is associated with the augmented PHY.

A. FER performance

The PER and HER performances were investigated using Monte Carlo simulations, which were each continued until we observed 1 000 erroneous payloads or headers, as appropriate. In common with [1, Figure E.2], we considered transmissions over line-of-sight Additive White Gaussian Noise (AWGN) channels. Here, a range of channel Signal to Noise Ratios (SNRs) per payload chip E_{pc}/N_0 and per header chip E_{hc}/N_0 [1, Section E.5.5.1] were considered. Note that these E_{pc}/N_0 and E_{hc}/N_0 values were selected to be in excess of the schemes' channel capacity bound of -10.26 dB [3]. Our results are plotted in Figure 3.

Observe in Figure 3 that the standard PHY without augmentation achieves lower PERs when shorter payloads are

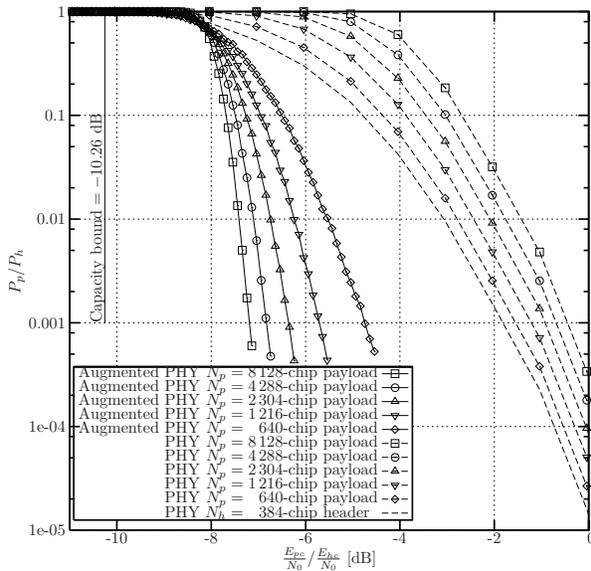


Fig. 3. PER and HER performance.

employed, since these comprise less chips that may be corrupted. By contrast, the augmented PHY achieves lower PERs when longer payloads are employed, since these employ longer DRP interleavers, as discussed in Section II-A. This beneficial effect outweighs the above-mentioned increased corruption probabilities that are associated with longer payloads. Note that the PER results plotted in Figure 3 for the augmented PHY using DRP interleavers were found to be indistinguishable from those obtained using randomly designed interleavers [21].

Let us now compare the PER performances of the augmented and standard PHYs in scenarios where they convey an equal number of N_p payload chips in the presence of noise having the same power spectral density of N_0 . As predicted in Section II, the augmented PHY always achieves a lower PER P_p than the standard PHY without augmentation, regardless of how much transmit energy per payload chip E_{pc} is employed, as shown in Figure 3. Hence, the amount of sensor node energy that is consumed by retransmitting erroneously received data is significantly reduced, when employing the augmented PHY instead of the standard PHY, as discussed in Section I.

In an alternative interpretation of Figure 3, the augmented PHY can always achieve a particular PER P_p using a lower transmit energy per payload chip E_{pc} than the standard PHY. Indeed, the discrepancies between the channel capacity bound of -10.26 dB and the augmented PHY's PER plots are approximately equal to half of the discrepancies associated with the standard PHY without augmentation, as shown in Figure 3. Table I lists the E_{pc} gains that may be achieved by transmitting the PHY payload using the augmented PHY instead of the PHY without augmentation at a PER of $P_p = 0.001$.

N_p	E_{pc} gain for $P_p = 0.001$	E_c gain for $P = 0.001$
640	3.22 dB	1.81 dB
1 216	4.50 dB	2.48 dB
2 304	5.42 dB	3.09 dB
4 288	6.16 dB	3.57 dB
8 128	6.75 dB	3.99 dB

TABLE I

TRANSMISSION ENERGY GAIN ACHIEVED BY TRANSMITTING THE PHY PAYLOAD USING THE AUGMENTED PHY INSTEAD OF THAT WITHOUT AUGMENTATION.

Note that the PERs plotted in Figure 3 assume that the PHY headers are received without error. This is because the PHY headers convey the synchronisation sequence and the PHY payload length, both of which must be recovered without error before iterative decoding may commence in the proposed augmentation to the PHY. In the more realistic scenario where error-free PHY headers cannot be guaranteed, the FER P can be obtained as

$$P = P_h + P_p(1 - P_h). \quad (1)$$

In this case, the overall transmit energy per chip E_c is given by

$$E_c = \frac{N_p E_{pc} + N_h E_{hc}}{N_p + N_h}. \quad (2)$$

Here, the values of E_{pc} and E_{hc} may be specially selected in order to minimize the overall transmit energy per chip E_c that is required to obtain a particular FER P . When the PHY payload is conveyed using the standard PHY without augmentation, the transmit energy per payload chip E_{pc} should equal the chip energy E_{hc} that is employed to transmit the headers, since these are also conveyed using the standard PHY. By contrast, when the PHY payload is conveyed using the augmented PHY, the overall transmit energy per chip E_c is minimized by employing different values of E_{pc} and E_{hc} . We therefore conducted a search to find the optimal values of E_{pc} and E_{hc} that yield FERs of $P = 0.001$ for each PHY payload length N_p considered. For the case where $N_p = 640$, the optimal value of E_{hc} was found to be 2.78 dB higher than the optimal value of E_{pc} . This discrepancy was found to increase with the PHY payload length, equalling 3.89 dB when $N_p = 8128$. Table I lists the gains in the minimized overall transmit energies per chip E_c that may be achieved by employing the augmented PHY instead of the standard PHY to achieve an FER of $P = 0.001$.

Note that the transmit energy gain that can be realized in practice is limited by the requirement for the signal to be readily differentiated from noise by the Carrier Sense Multiple Access and Collision Avoidance (CSMA-CA) mechanism [1, Section 7.5.1.4] of all other transmitters in the network. More specifically, if a low-energy transmission is not detected by the CSMA-CA mechanism of a transmitter that is ready to send, then it will generate a significant amount of interference for the low-energy transmission. This may be avoided by employing slotted CSMA-CA [1, Section 7.5.1.4], in which

transmissions can only commence at regular intervals. As a result, the CSMA-CA mechanism of a potentially interfering transmitter will be invoked at the start of the augmented transmissions, during the synchronisation and PHY headers, which are transmitted with relatively high energies using the standard PHY without augmentation, as described above. If we assume that this transmit energy allows the synchronisation and PHY headers to be differentiated from noise by the CSMA-CA mechanism of the other transmitter that is ready to transmit, then interference will be avoided.

B. Encoding complexity and overall transmitter energy consumption

Let us now estimate how much energy would be consumed in practice by the additional interleaving and rate-1 encoding processes that are invoked by the augmented PHY's transmitter. This can then be used to offset the transmit energy savings detailed in Section III-A.

One possible implementation of our augmented PHY's transmitter would resemble the Chipcon CC2430 [22], but with the addition of a module dedicated to performing interleaving and rate-1 encoding. As described in Section II-B, we can conservatively estimate that this module would require $4N_p$ clock cycles to process an N_p -chip data payload. The duration of the processing is therefore $t = 4N_p/f_{pr}$, where we assume that the clock speed is $f_{pr} = 32$ MHz, which equals that of the Chipcon CC2430 [22]. The energy consumed by the dedicated module is given by $E_{pr} = I_{pr}Vt = 4I_{pr}VN_p/f_{pr}$, where we assume the same supply voltage of $V = 3$ V as the Chipcon CC2430 and we conservatively assume $I_{pr} = 12.3$ mA, which equals the peak current consumption of the 8051 microcontroller on the Chipcon CC2430 [22, Table 4]. The energy E_{pr}^{aug} consumed when processing data payloads comprising various numbers of chips N_p is provided in Table II.

N_p	PHY	Augmented PHY		Gain
	E_{tx}^{std}	E_{tx}^{aug}	E_{pr}^{aug}	$E_{tx}^{std} - (E_{pr}^{aug} + E_{tx}^{aug})$
640	31.1 μ J	22.7 μ J	3.0 μ J	5.4 μ J
1216	59.1 μ J	41.8 μ J	5.6 μ J	11.7 μ J
2304	112.0 μ J	77.4 μ J	10.6 μ J	24.0 μ J
4288	208.4 μ J	141.5 μ J	19.8 μ J	47.1 μ J
8128	395.0 μ J	264.6 μ J	37.5 μ J	91.9 μ J

TABLE II
TOTAL ENERGY GAIN ACHIEVED BY TRANSMITTING THE PHY PAYLOAD USING THE AUGMENTED PHY INSTEAD OF THAT WITHOUT AUGMENTATION.

The energy consumed during transmission is given by $E_{tx} = I_{tx}VN_p/f_{tx}$, where the IEEE 802.15.4 transmission rate is $f_{tx} = 2 \times 10^6$ chips per second [1]. As may be expected, the current I_{tx} consumed during the transmission of a data payload depends on the particular transmit energy per chip E_{pc} employed. In its maximum transmit power mode of 0.6 dBm, the Chipcon CC2430 consumes $I = 32.4$ mA [22, Table 45]. At this transmit power, the amount of energy E_{tx}^{std} consumed by the standard PHY without augmentation is provided in Table II for various values of N_p . However, invoking the

augmented PHY facilitates reductions of 3.22 – 6.75 dB in the required transmit energy per chip E_{pc} , as shown in Table I. Corresponding reductions from the Chipcon CC2430's maximum transmit power of 0.6 dBm allow us to lower its current consumption to 21.7 – 23.7 mA [22, Table 45]. Table II provides the amount of energy E_{tx}^{aug} consumed by the augmented PHY for various values of N_p , as well as the overall associated energy savings $E_{tx}^{std} - (E_{pr}^{aug} + E_{tx}^{aug})$, both of which are about 20% of E_{tx}^{std} . Note that even higher energy savings may be expected in practice, owing to the conservative assumptions made in this analysis.

C. Decoding complexity

Let us now quantify the decoding complexity increase associated with the augmented PHY. As described in Section II-C, the iterative decoding process of the augmented PHY can be completed using only fixed-point ACS operations. During our Monte Carlo simulations, we recorded the number of ACS operations per payload chip that was required to reach iterative decoding convergence, where no further PER performance improvement could be obtained by undertaking additional decoding iterations. These ACS operation counts are plotted as a function of E_{pc}/N_0 in Figure 4 for each considered number of payload chips N_p .

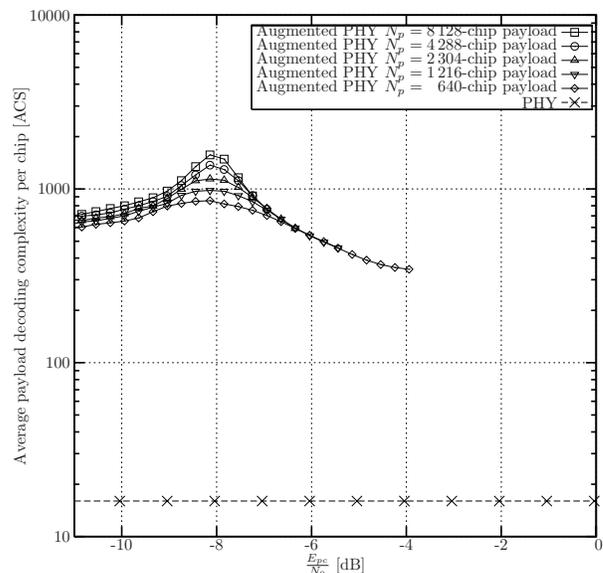


Fig. 4. PHY payload decoding complexity.

As shown in Figure 4, the decoding complexity of the augmented PHY peaks at about $E_{pc}/N_0 = -8$ dB, which is also the threshold between the regions of high and low PERs in Figure 3. At E_{pc}/N_0 values below -8 dB, the iterative decoding process is barely able to achieve any error correction and hence quickly converges, yielding a relatively low decoding complexity. By contrast, the iterative decoding process achieves substantial error correction improvements for E_{pc}/N_0

values above -8 dB, rapidly converging to a low PER and, again, yielding a relatively low decoding complexity. Hence, the decoding complexity peak observed may be explained by the facilitated, but gradual, error correcting progress that can be achieved at $E_{pc}/N_0 = -8$ dB.

In contrast to the augmented PHY, the decoding complexity per payload chip of the standard PHY without augmentation is independent of both the E_{pc}/N_0 value and the number N_p of payload chips considered. This is because the standard PHY employs only a single decoding iteration rather than a variable number of iterations. As shown in Figure 4, the peak decoding complexity of the augmented PHY is 100 times higher than that of the PHY without augmentation. However, this peak complexity is not relevant, since the PER P_p at the corresponding E_{pc}/N_0 value of -8 dB is high, as shown in Figure 3. Indeed, the decoding complexity of the augmented PHY is no more than 56 times higher than that of the standard PHY for E_{pc}/N_0 values in excess of -7.2 dB, where PERs of $P_p < 10^{-3}$ can be achieved. While significant, we consider this increased complexity to be reasonable, since it is only incurred by the central WSN node, which we assume to have abundant energy resources, as described in Section I. Also note that the decoding complexity may be significantly reduced by halting the iterative decoding process, before convergence is achieved. Note however, that the resultant PER may be affected if the iterative decoding process is halted too early.

IV. CONCLUSIONS

In this paper, we have characterized an augmentation of the IEEE 802.15.4 [1] 2450 MHz physical layer. This augmentation was shown to facilitate a significantly improved FER performance and/or a significant transmit energy gain of $1.81 - 3.99$ dB, at the cost of a decoding complexity increase of up to 56 times (as well as a marginally higher encoding complexity). As discussed in Section I, this trade off is particularly desirable in WSNs, where the sensor nodes have scarce energy resources (obtained through energy harvesting for example) and where all data frames are routed via a central node having an abundant energy resource (such as the mains). Indeed, we demonstrated that our approach can reduce the overall energy consumption of the transmitting sensor nodes by more than 20% in practice.

REFERENCES

- [1] *Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs)*, IEEE Std. 802.15.4, September 2006.
- [2] L. Hanzo, L.-L. Yang, E.-L. Kuan, and K. Yen, *Single and Multi-Carrier DS-SS-CDMA: Multi-User Detection, Space-Time Spreading, Synchronisation and Standards*. John Wiley and Sons, 2003.
- [3] L. Hanzo, S. X. Ng, T. Keller, and W. Webb, *Quadrature Amplitude Modulation*. Chichester, UK: Wiley, 2004.
- [4] J. Hagenauer, E. Offer, and L. Papke, "Iterative decoding of binary block and convolutional codes," *IEEE Transactions on Information Theory*, vol. 42, no. 2, pp. 429–445, March 1996.
- [5] R. Herzog, J. Hagenauer, and A. Schmidbauer, "Soft-in/soft-out Hadamard despreader for iterative decoding in the IS-95(A) system," in *Proceedings of the IEEE Vehicular Technology Conference*, vol. 2, Phoenix, AZ, USA, May 1997, pp. 1219–1222.
- [6] S. Benedetto and G. Montorsi, "Serial concatenation of block and convolutional codes," *Electronics Letters*, vol. 32, no. 10, pp. 887–888, May 1996.
- [7] D. Divsalar, S. Dolinar, and F. Pollara, "Serial concatenated trellis coded modulation with rate-1 inner code," in *Proceedings of the IEEE Global Telecommunications Conference*, vol. 2, San Francisco, CA, USA, November 2000, pp. 777–782.
- [8] S. Benedetto and G. Montorsi, "Iterative decoding of serially concatenated convolutional codes," *Electronics Letters*, vol. 32, no. 13, pp. 1186–1188, June 1996.
- [9] M. E. Magana, D. H. Buck, and J. Speidel, "Improvement of data rate performance of wireless local area networks through the use of orthogonal code division multiplexing (OCDM), multipath diversity, and parallel concatenated channel encoding," *Wireless Personal Communications*, vol. 25, no. 4, pp. 293–305, July 2003.
- [10] M. Bykowski, "Performance of turbo codes with SOVA algorithm in DSSS over channels with non-white additive Gaussian noise," in *Proceedings of the Advanced Industrial Conference on Telecommunications*, Lisbon, Portugal, July 2005, pp. 226–229.
- [11] H. A. Cirpan and M. K. Tsatsanis, "Chip interleaving in direct sequence CDMA systems," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 5, Munich, Germany, April 1997, pp. 3877–3880.
- [12] R. H. Mahadevappa and J. G. Proakis, "Mitigating multiple access interference and intersymbol interference in uncoded CDMA systems with chip-level interleaving," *IEEE Transactions on Wireless Communications*, vol. 1, no. 4, pp. 781–792, October 2002.
- [13] B. Vucetic, Y. Li, L. C. Perez, and F. Jiang, "Recent advances in turbo code design and theory," *Proceedings of the IEEE*, vol. 95, no. 6, pp. 1323–1344, June 2007.
- [14] S. Crozier and P. Guinand, "High-performance low-memory interleaver banks for turbo-codes," in *Proceedings of the IEEE Vehicular Technology Conference*, vol. 4, Atlantic City, NJ, USA, October 2001, pp. 2394–2398.
- [15] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, 1989.
- [16] L. Hanzo, T. H. Liew, and B. L. Yeap, *Turbo Coding, Turbo Equalisation and Space Time Coding for Transmission over Wireless Channels*. Chichester, UK: Wiley, 2002.
- [17] I. Land, P. Hoeher, and S. Gligorević, "Computation of symbol-wise mutual information in transmission systems with log APP decoders and application to EXIT charts," in *Proceedings of the International ITG Conference on Source and Channel Coding*, Erlangen, Germany, January 2004, pp. 195–202.
- [18] S. ten Brink, "Convergence of iterative decoding," *Electronics Letters*, vol. 35, no. 10, pp. 806–808, May 1999.
- [19] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv, "Optimal decoding of linear codes for minimizing symbol error rate (Corresp.)," *IEEE Transactions on Information Theory*, vol. 20, no. 2, pp. 284–287, March 1974.
- [20] G. D. Forney, "Review of random tree codes," NASA Ames Research Center, Moffett Field, CA, USA, Tech. Rep. NASA CR73176, December 1967.
- [21] L. Dinoui and S. Benedetto, "Design of fast-prunable S-random interleavers," *IEEE Transactions on Wireless Communications*, vol. 4, no. 5, pp. 2540–2548, September 2005.
- [22] *A True System-on-Chip Solution for 2.4 GHz IEEE 802.15.4 / Zig-Bee(TM) Datasheet*, Chipcon, June 2007.