Multi-Rate Selection in ZigBee

Linghe Kong, *Senior Member, IEEE*, Yifeng Cao, Liang He, *Senior Member, IEEE*, Guihai Chen, Min-You Wu, *Senior Member, IEEE*, and Tian He, *Fellow, IEEE*

***Abstract*— ZigBee is a widely used wireless technology in low-power and short-range scenarios such as the Internet of Things, sensor networks, and industrial wireless networks. However, the traditional ZigBee supports only one data rate, 250 Kbps, which thoroughly limits ZigBee's efficiency in dynamic wireless channels. In this paper, we propose Mrs. Z, a novel physical layer design to enable multi-rate selection in ZigBee with lightweight modification on the legacy ZigBee modules. The key idea is to change the single spectrum spreading length to multiple ones. Correspondingly, to support the rate adaptation to the channel variations, we propose a bit-error-based rate selection scheme, which predicts BER by leveraging the physical properties of ZigBee to calculate the confidence for each symbol in transmission. Then, the receiver selects the rate based on the negative impact on throughput incurred by bit errors and gives feedback to the transceiver. We implement Mrs. Z on USRPs and evaluate its performance in different scenarios. Experiment results demonstrate that Mrs. Z achieves about 1.15, 1.2, and 1.8 x average throughput compared to the classic smart pilot, softrate, and the traditional ZigBee.**

***Index Terms*— Multi-rate selection, rate adaptation, ZigBee.**

I. Introduction

ZigBee, a low-cost, low-rate, and low-power communication technology based on the IEEE 802.15.4 standard [2], has been widely adopted in various Internet of Things (IoT) applications, such as smart home [3] and smart-grid systems [4]. Such wide deployments of ZigBee are currently faced with three challenges: (i) requirements for the higher data rate in data-incentive IoTs, (ii) resistance to complex and dynamic channels and (iii) adaptivity to limited wireless link resources, all of which impose ever-increasing requirements on its throughput. Techniques including collision avoidance [5] and collision resolution [6] improve the effective throughput by reducing packet collisions. For example, ZigZag [7] reduces communication collisions by separating packets with time offsets, and mZig [8] recovers the transmitted information from collided packets by decomposing them chip-by-chip. However, these methods are still strictly limited by ZigBee's only data rate, i.e. 250Kbps.

Multi-rate adaptation — i.e, adaptively adjusting the communication data rate based on the real-time link quality — is another orthogonal dimension to improve the throughput, whose effectiveness has been widely proved in WiFi-based communications, but is still defectively covered in Zigbee. The 802.11 standard, the core of WiFi, is able to offer varying data rates of 6∼54Mbps [9], by adjusting its modulation schemes (e.g., BPSK, QPSK, 16-QAM and 64-QAM in 802.11n) and coding rates (e.g., 1/2, 2/3, 3/4 and 5/6) at the transceiver. To enhance WiFi's throughput against channel variations, researchers propose various rate selection schemes. Typically, rate adaptation can be classified into two categories: MAC-layer based and physical-layer based schemes. MAC-layer based schemes, e.g, SampleRate [10] and RRAA [11], adapt rates mainly based on packet losses. They are easy to implement but insufficient in responsiveness. Thus, physical-layer based schemes such as SoftRate [12] and CHARM [13] are proposed, which exploit physical hints including SNR or BER for fine-grained calculation on channel quality. Physical-layer based rate adaptation outperforms previous schemes both in accuracy and responsiveness. Based on the observations on 802.11, we are motivated to explore the opportunity to enable multi-rate selection with specially-designed rate selection scheme in ZigBee, extending ZigBee's intrinsic throughput bound.

However, there are two challenges to make rate adaption in ZigBee. First, ZigBee supports only one modulation scheme in the physical layer, leading to the fixed data rate, i.e., 250Kbps with the *Offset-Quadrature Phase Shift Keying* (O-QPSK) physical layer at 2.4GHz [2], impeding the deployment of rate adaptation techniques thereon. One approach is adding new modulation schemes to ZigBee but this clearly brings in huge rewriting cost in hardware. Lightweight modification is comparatively more fascinated. Second, we observe that current rate selection schemes, mainly in WiFi, cannot achieve as good performance in ZigBee. For example, SoftRate calculates incorrect bits with *Log-Likelihood Ratios* (LLR) directly in WiFi. However, it is infeasible in ZigBee. A ZigBee transceiver will spread bits to chip sequences in the physical layer, covering up the characteristics of original bits. As a result, using LLRs to estimate BER suffers severe deviation.

Manuscript received November 28, 2018; revised March 25, 2019; accepted March 25, 2019; approved by IEEE/ACM Transactions on Networking Editor C. F. Chiasserini. Date of publication May 3, 2019; date of current version June 14, 2019. This work was supported in part by the National Key Research and Development Program of China under Grant 2018YFB1004703, in part by the NSFC under Grant 61672349, Grant 61672353, and Grant 61672349, and in part by the Shanghai Key Laboratory of Scalable Computing and Systems, Shanghai Jiao Tong University. This paper extends our previous work published in IEEE ICNP17 [1].

(Corresponding author: Linghe Kong.)

L. Kong, Y. Cao, G. Chen, and M.-Y. Wu are with the Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: linghe.kong@sjtu.edu.cn; yyshxyh2013@sjtu.edu.cn; chen-gh@sjtu.edu.cn; mwu@sjtu.edu.cn).

L. He is with the Department of Computer Science and Engineering, University of Colorado Denver, Denver, CO 80204 USA (e-mail: liang.he@ucdenver.edu).

T. He is with the Department of Computer Science and Engineering, University of Minnesota, MN 55455 USA (e-mail: tianhe@cs.umn.edu).

Digital Object Identifier 10.1109/TNET.2019.2913014

To close the gap, we propose *multi-rate-selection in ZigBee* (Mrs.Z), a novel physical layer design that enables rate adaptation in ZigBee with lightweight modification on commodity modules, improving its throughput while ensuring reliability and scalability. Mrs.Z is designed and implemented on the top of *direct-sequence-spread-spectrum* (DSSS), a widely used technique in ZigBee to reduce the negative impacts of interferences [2]. Specifically, DSSS improves communication reliability by encoding the original message to a longer chip sequence generated with pseudo-random code. This, however, is at the cost of reduced throughput. Mrs.Z enhances DSSS by adaptively adjusting its coding rates (and thus the DSSS chip sequences) based on real-time link quality, with the objective to maximize the throughput. Such adaptation, in turn, is determined based on an online predicted *bit error rate* (BER). To estimate BER, Mrs.Z calculates the confidence of each chip and combines them to obtain the confidence of each symbol through maximum likelihood methods. Because the symbol confidence is almost linearly dependent on BER, Mrs.Z can finally predict BER accurately from the raw samples in the receiver.

We implement and evaluate Mrs.Z on USRP N210 testbed, with both static and mobile channels. The results show Mrs.Z improves the throughput by 15%, 20% compared to Smart-Pilot [14], SoftRate [12], two existing rate selection schemes originally developed for WiFi, and by 80% when compared with the traditional ZigBee. In high-SNR links, the throughput can even better up to 1.4x, 1.4x and 2.9x of Smart-Pilot, SoftRate and traditional ZigBee. This paper makes the following contributions:

- We propose Mrs.Z to enable multi-rate transmission in ZigBee, uncovering the limitation of the single and fixed data rate on ZigBee's throughput;
- We present a rate selection scheme, a physical layer enhancement of standard ZigBee, which adapts the data rate effectively and efficiently;
- We implement and evaluate Mrs.Z on USRPs, showing 1.15x, 1.2x and 1.8x performance on average compared to Smart-Pilot, SoftRate and traditional ZigBee, respectively.

The paper is organized as follows. Section II reviews the related literature. Section III introduces the background and our motivation. The design and implementation of Mrs.Z are presented in Section IV and V, respectively. Mrs.Z is evaluated in Section VI, and the paper concludes in Section VII.

II. Related Works

Rate adaptation was a well-addressed topic in wireless communication in the past years but is still not explored in ZigBee. Typically, rate adaptation schemes can be divided into MAC-layer based rate selection and physical-layer based rate selection.

A. MAC-Layer Schemes

MAC-layer based rate selection schemes, including RRAA [11], WOOF [15], SampleRate [10], etc., make rate selection with *packet loss rate* (PLR). ARF [16] is one of the classic rate selection algorithm based on the packet loss. CARA enhances ARF by distinguishing collisions from channel fading when packet loss occurs. Note that although collisions incur packet loss, the rate should not be decreased because the collision does not reflect the properties of channels. To make a distinction, CARA [17] leverages *Request-to-Send* (RTS) and *Clear-to-Send* (CTS) for each packet loss to reduce unintended rate decrease. RRAA extends the above PLR-based schemes by using frame loss information gathered over tens of frames to adapt the rate. Another important improvement in RRAA is that it reduces the overhead incurred through RTS/CTS.

WOOF removes the overhead of RTS/CTS by using channel busy time to monitor the network load. With higher channel busy time, a transmission failure is more likely to be caused by a collision. SampleRate, which has been deployed on Atheros cards, takes a different method. It makes a prediction on transmission time by frequently sampling with different transmission rates and tries to minimize the average transmission time. However, SampleRate does not differentiate between collision and noise.

Generally, while MAC-layer based rate selection schemes are easy to implement, they lack responsiveness to the channel variations because the receiver usually needs multiple receptions for one rate selection.

B. Physical-Layer Schemes

Different from MAC-layer schemes, physical-layer based rate selection exploits physical hints to select an appropriate rate. SGRA [18] predicts frame delivery ratio directly from SNR. It calculates SNR from *received signal strength indication* (RSSI). However, RSSI is known to be coarse and the SNR-BER relation varies over real wireless channels. Similarly, CHARM is a scheme leveraging reciprocity of wireless channel to estimate average SNR at the receiver, thereby picks a rate based on SNR. The problem lies in that the reciprocity is not always reliable in testbeds, leading to a sub-optimal rate selection. ESNR [19] enhances SNR measurement by considering frequency diversity and calibrates SNR with CSI. However, ESNR still fails to select the optimal rate due to the limited physical information.

To reduce the impact of SNR variation, SoftRate predicts BER with LLR for rate selection. With physical layer hints, it estimates per-frame BER as a feedback to the transceiver, where a new rate can be picked for the next transmission. To accurately distinguish collisions from channel fading, it adds a postamble in the end of each frame to detect a collision with high likelihood. AccuRate shows there is still some improvement room in SoftRate. AccuRate [20] estimates the channel condition by observing the dispersion of the received symbols. Then, it can select the highest rate whose dispersion is below the threshold, which is more fine-grained.

For further improvements in throughput, some research seeks for more hints to calibrate channel quality [14], [21]–[23]. H-RCA further takes machine learning approaches, which employs Bayesian analysis for each rate-increase trial, ensuring the rate increase will not lead to a poorer performance. Meanwhile, TXOP [24] technique is applied to accurately distinguish collision loss. SmartPilot correlates multiple physical hints to improve the accuracy of obtaining the optimal rate. It exploits potential pilots across multiple layers to track the channel variance. Such pilots come from two cases. First, decoded bits with high confidence are termed as soft pilots. Second, headers whose values are fixed can also improve the channel estimation, which are termed as hard pilots.

However, all of the above rate selection schemes fail to achieve optimal performance in ZigBee because of the differences in the physical layer. According to the standard, ZigBee spreads bit stream to chip sequences to enhance the reliability in transmission, where all of the above schemes do not give a proper approach to estimate BER from chip errors, which can be obtained by the receiver. In addition, the spreading scheme is fixed in ZigBee, thus the ground-truth candidate chip sequences are known a priori. We believe this feature can be explored to make a specially-designed rate selection scheme for ZigBee.

C. Other Schemes

In the recent years, new scenarios or metrics are considered into rate selection schemes, making the rate adaptation more complicated. As MIMO is supported in 802.11n, MiRA [25] zigzags between intra- and inter-mode rate adaptations so that the rate selection can still work smoothly in diversity-oriented, single-stream mode or spatial multiplexing driven, double-stream mode. TurboRate [26] enables rate selection for multi-user LANs. TurboRate picks rates based on the SNR in single-in-single-out scenario and the direction of the transceiver's signal received by the access point. EERA [27] considers energy consumption in rate selection. It achieves a trade-off between the energy efficiency and higher throughput.

III. PRELIMINARY

Here we briefly review the conventional design of ZigBee's physical layer and the potential opportunity to enable multi-rate transmission in ZigBee.

ZigBee operates in four different bands of 780MHz, 868MHz, 915MHz and 2.4GHz. In this paper, we mainly focus on ZigBee in the worldwide frequency band of 2.4GHz, with a 250Kbps data rate.

In ZigBee, the *physical layer* (PHY) first encapsulates data from the MAC layer to generate the *PHY protocol data unit* (PPDU), which contains (i) the *synchronization header* (SHR) field as a header, (ii) the *physical header* (PHR) field specifying the frame length, and (iii) the payload carrying the to-be-transmitted data.

A. Transceiver

The *transceiver* (TX) sends the PPDU to the receiver with the following steps: bit-to-symbol mapping, DSSS, O-QPSK modulation, and pulse shaping, as illustrated in Fig. 1.

In the bit-to-symbol mapping, the binary stream in PPDU is encoded to a symbol stream. For each octet in the binary stream, its four *least significant bits* (LSB) (b_0, b_1, b_2, b_3) are mapped to one 4-bit symbol, and four *most significant bits* (MSB) (b_4, b_5, b_6, b_7) are mapped to the next symbol. This mapping applies to the entire PPDU and the output symbol stream is sent into DSSS.

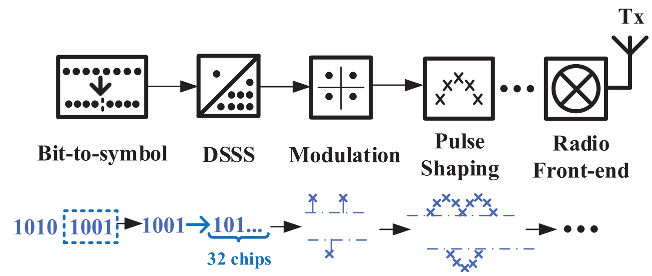

Fig. 1. Block diagram of ZigBee's transceiver in physical layer.

In the DSSS phase, each symbol is spread to a 32-chip pseudo-random noise sequence, where a chip is the smallest information-carrying unit in ZigBee. DSSS is used to enhance the transmission reliability against potential interferences.

In the O-QPSK modulation phase, the chip sequence from DSSS module is modulated as follows: even-indexed chips are modulated onto the *in-phase* (I) carrier, and odd-indexed chips are modulated onto the *quadrature-phase* (Q) carrier. Q-phase chips are delayed by half-chip time with respect to I-phase chips to form an offset.

In the pulse shaping phase, the chip sequence is shaped into half-sine pulse and then transmitted after digital-analog conversion. The half-sine pulse is given by:

$$S(t) = \begin{cases} \sin(\pi \dfrac{t}{2T}), & 0 \leq t \leq 2T_c; \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

B. Receiver

The receiver (RX) down-converts the ZigBee waveforms to the baseband and digitalizes them to complex samples $s(n)$ with ADC. Corresponding to the OQPSK, quadrature demodulation is used to recover chips from samples. Specifically, the RX obtains the phase shift $\triangle\phi$ by calculating $s(n) \times s^*(n-1)$ ($s^*(n-1)$ is the conjugate of $s(n-1)$). With no noise and channel fading, the phase shift $\triangle\phi$ should be a constant in one-symbol duration. If $\triangle\phi > 0$, the RX outputs the chip value '1' and otherwise outputs '0'.

When decoding symbols from the received chip sequence, the receiver compares the chip sequence with each element in the symbol-chip mapping table, a table maintaining the corresponding chip sequence for each symbol. The receiver despreads the received chip sequence to the candidate symbol — the symbol with the least different chips when compared to the received chip sequence — if the number of such different chips is smaller than a pre-defined decoding threshold d_{th}; Otherwise, it concludes the transmission as failure and dumps the received chip sequence.

As specified in the standard, ZigBee has a bandwidth of 2MHz [2]. DSSS, however, spreads each 4-bit symbol to a 32-bit chip sequence to ensure the transmission reliability, incurring redundancy in the transmitted chips and thus waste of bandwidth. As a result, only a 250Kbps data rate, or one-eighth of the available bandwidth, is achieved

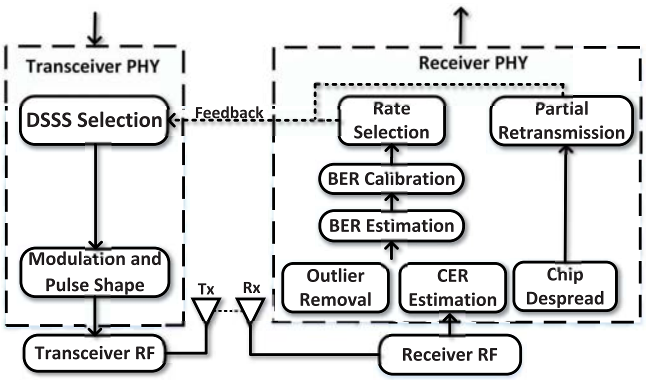

Fig. 2. An overview of Mrs.Z.

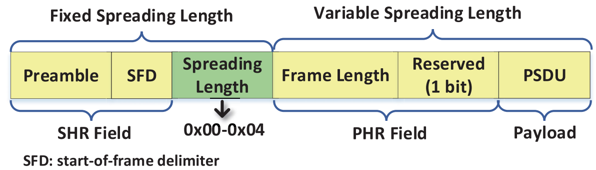

Fig. 3. The packet format in Mrs.Z.

Symbol	Chip Sequence
...	...
0001	0100111110001001
...	...
1111	1010110111000001

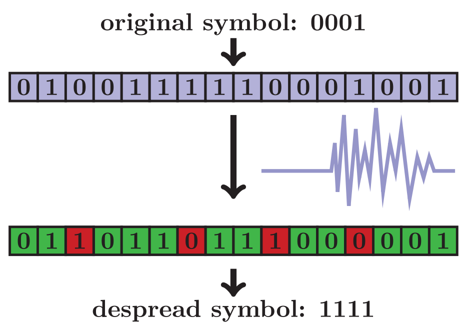

Fig. 4. In this case, the receiver incorrectly despreads the chip sequence of 0001 to another symbol 1111.

in practice, rendering it possible to improve the throughput by reducing transmission redundancy. On the other hand, in an extremely-noisy channel, even 4-32 spreading is unreliable, which requires for a longer spreading length.

IV. Multi-Rate Transmission of ZigBee

We present the design of Mrs.Z in this section. We first explain how we achieve the multi-rate transmission with current ZigBee modules, and then we introduce the online rate selection scheme to determine the optimal transmission rate.

Fig. 2 gives an overview of Mrs.Z. At the transceiver, an extra module for selection is plugged in to adapt the rate accordingly. At the receiver, a packet containing the data rate and retransmission information is generated as a feedback to the transceiver.

A. Multi-Rate Transmission

To remedy the weakness of fixed-length DSSS in traditional ZigBee, Mrs.Z uses five different chip lengths for adaptive rate selection to make full of the bandwidth. Before each packet transmission, the TX selects the proper chip length to maximize the throughput according to the feedback from the RX. Specifically, Mrs.Z supports 4-to-4 mapping (i.e., no redundancy), 4-to-8 mapping, 4-to-16 mapping, 4-to-32 mapping, and 4-to-64 mapping. The low-redundancy spreadings including 4-to-4, 4-to-8 and 4-to-16 mappings, can be used to achieve higher throughput in high-SNR channels, while the 4-to-32 and 4-to-64 mappings will be used to ensure reliability in low-SNR channels.

Codeword Design: The encoding of Mrs.Z is designed under the rules in legacy ZigBee. Specifically, the codeword design satisfies the following two properties:

- Odd-even symmetry. The number of chip '1' is equal to chip '0' in each chip sequence.
- Cyclicity. Except symbol '0' and '8', each chip sequence is obtained by shifting the last chip sequence right for n bits.

As special illustration, when we use 4-4 spreading, the TX actually uses 5 chips to encode a symbol. We do so because in O-QPSK demodulation, when the RX resolves a packet as illustrated in Section III-B, the first chip is always skipped. Thus we need 5 chips to explicitly differentiate all the symbols.

Transceiver Operations: Mrs.Z uses an extra byte in PPDU to represent the spreading length L_s ($L_s \in \{4, 8, 16, 32, 64\}$), as shown in Fig. 3 (Four bits are enough but here we use one byte to meet the industrial standard). To ensure the critical information in a packet is transmitted correctly, Mrs.Z divides the packet into fixed-length and variable-length parts. Specifically, the transceiver always spreads the SHR field and and the extra bits to 64 chips, ensuring their reliability in RX. Other bits, e.g., the PHR field and payload, are spread according to the RX's feedback in previous transmission. The TX will receive a feedback to update its DSSS length from RX after each transmission. A retransmission is triggered if no feedback comes in expiration.

Receiver Operations: The RX is responsible for decoding and rate adaptation. Upon detecting a packet, the RX uses 64-to-4 despreading to resolve L_s, which is then used to recover the other fields. With L_s chips as a segment unit, the RX compares the decoded chips with the standard chips by referring to the DSSS table and decodes the chip sequence to the corresponding symbol. On the other hand, the RX operates on raw samples to evaluate channel conditions. Then the RX sends a feedback packet, informing the TX of the new spreading length and the packet-transmission state. The L_s for feedback data is set as 64 chips. The extra overhead of feedback is tolerable because the payload length is comparatively negligible in the transmission.

Effectiveness Evaluation: Bit errors occur when the decoded sequence has the smallest hamming distance with another symbol instead of the ground truth. As shown in Fig. 4, the TX encodes original symbol '0001' to 16-chip

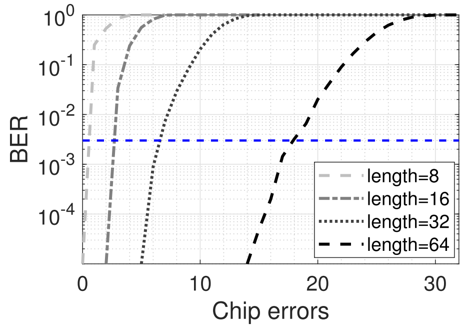

Fig. 5. The effect of chip errors in the sequence on BER.

sequence '0100111110001001'. Because of the existence of noise and attenuation, the RX obtained the sequence '0110110111000001', with four chip errors. By the principle of decoding, this chip sequence will be finally decoded into symbol '1111' with less different chips compared to the ground truth '0001', which incurs bit errors.

Based on the effects of chip errors on bit errors, we make simulations to theoretically evaluate whether the design of multi-rate transmission is effective. Intuitively, to ensure a longer sequence is more robust than a shorter one, it has to obtain lower BER in transmission. Specifically, say, the mean BER is b_l when the spreading length is l with n chip errors in the sequence. Then the length is effective when BER $b_{2l} < b_l$ with $2n$ chip errors in the sequence. The effect of chip errors on BER is shown in Fig. 5. Note that we do not plot the case when $l = 4$ because any chip error will produce a bit error in this case. We also plot the baseline when BER is 3×10^{-3} that is usually used as an industrial standard [8]. From a general sense, to keep BER in 3×10^{-3}, the tolerable chip errors are $0, 2, 6, 18$ when $l = 8, 16, 32, 64$, respectively, which demonstrates that $l = 16, 32, 64$ are *effective lengths*. When $l = 8$, no chip error is allowed to satisfy $b < 3 \times 10^{-3}$. Thus $l = 8$ will not be used in this sense, but it is still useful when some specials protocol are applied.

B. Rate Selection Scheme

Next we present the rate selection scheme of Mrs.Z.

Like all the other rate selection schemes, Mrs.Z aims at the throughput maximization. As fast responsiveness is a basic requirement to prevent channel variance, we implement our scheme in the physical layer.

The key is to tracking some physical hints reliably reflecting the throughput and predicting them accurately. We select BER as the metric. As proved in [12], BER is a good predictor of channel state and can be updated in a per-frame level. The difficulty lies in accurately estimating BER for each spreading length. For BER estimation in WiFi, SoftRate [12] uses LLRs to calculate overall *log*-confidence in the packet as BER and SmartPilot [14] calibrates *channel state information* (CSI) with both high-confidence bits in the frame body and hard information in the header. Unfortunately, they are imperfect when implemented in ZigBee directly. As we observe, BER estimation in ZigBee is faced with two difficulties:

(i) The effect of DSSS on BER is ambiguous. DSSS is an error-tolerant scheme. In spite of chip errors, BER is tightly related to the codeword design in DSSS. However, the effect of the codeword is almost unpredictable.

(ii) The preamble length of ZigBee packet is short (8 symbols). Using the preamble to estimate overall BER incurs severe estimation errors.

To surmount the above difficulties, Mrs.Z firstly calculates chip confidence one by one, and combines them to estimate the confidence of each symbol in DSSS. Here the chip confidence is obtained through the constellation. Specifically, the RX maps the dispersion of raw samples in one-chip time to the *error vector magnitude* in the constellation so that its confidence can be resolved. Then, the RX slides across the symbols in the DSSS table to obtain each symbol's confidence through the *maximum-likelihood* (ML) approach. Corresponding to the difficulties mentioned above, i) we have used the priori information in DSSS when we make ML-estimation to take the codeword into consideration; (ii) we estimate BER with the whole packet to improve the accuracy. Let the uncalibrated BER above be BER_u, the final BER is computed as follows:

$$\text{BER} = \alpha \text{BER}_\text{u}, \tag{2}$$

where α is a calibration factor depending on the spreading length l and CER. Note that α is required here in case of BER deviation in some scenarios, mainly in low-SNR channels. We will specify it in the section of calibration. With well-calibrated BER for the current length, the RX estimates BER of all lengths from the characterized relation between BER and SNR. Thus, the RX can calculate the effective throughput under each spreading length and select the proper rate accordingly.

The major steps are summarized as (i) BER estimation, (ii) BER calibration, (iii) cross-length BER estimation and (iv) rate selection.

Now, we explain them in detail.

1) BER Estimation: BER in ZigBee depends not only on each chip's confidence, but also the codeword design. In this section, we exploit a likelihood-based method to estimate BER in the RX. Specifically, the RX firstly calculates each chip's confidence and then uses them to further calculate how likely can each chip sequence be correctly decoded.

From the Section III, to decode a packet, the RX resolves the phase shift firstly and then maps it to '0' or '1'. Theoretically, the phase shift $\Delta\phi$ is a constant in one-chip duration. Actually, the value of $\Delta\phi$ is either $\tan(\frac{\pi}{n})$ or $-\tan(\frac{\pi}{n})$, where n depends on the sampling rate and bandwidth. However, with the existence of noise and channel fading, the calculated phase shift may disperse as shown in Fig. 6. Intuitively, the extent to which the phase shift disperses is a direct sign of BER. Next, Mrs.Z maps the phase shift with dispersion $\Delta\phi$ to a point (x, y) on the constellation. Specifically, suppose the samples in two consecutive chips are $s(i)$ $(i = 1, \ldots, N)$. An constellation point (x, y) is computed by:

$$\begin{cases} x = \dfrac{\frac{N}{2}\sum_{i=1}^{\frac{N}{2}} s(i)}{|\tan(\frac{\pi}{n})|}, \\ y = \dfrac{\frac{N}{2}\sum_{i=\frac{N}{2}+1}^{N} s(i)}{|\tan(\frac{\pi}{n})|}. \end{cases} \tag{3}$$

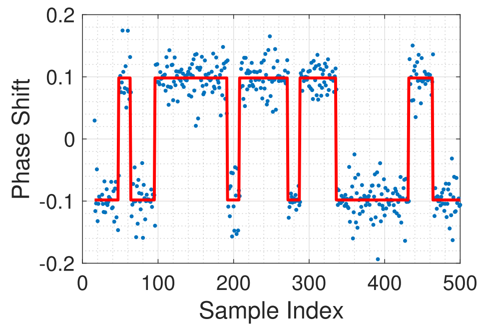

Fig. 6. Phase shifts disperse because of noise. The dispersion state can be described with a point in the constellation.

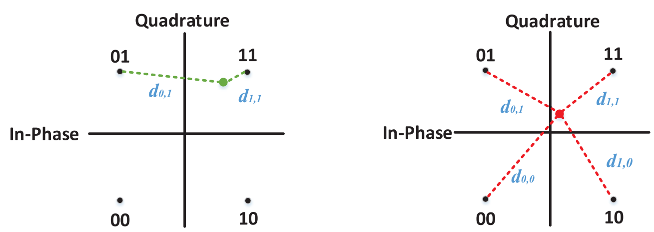

Fig. 7. High-confidence points have minor dispersion. In contrast, low-confidence points are dispersed apparently from the standard location.

Thus two chips are combined together as a point in the constellation as shown in Fig. 7, where its distance from the standard point reflects its confidence. Corresponding to O-QPSK, any constellation point is decoded based on their location, e.g., the point in the first quadrant corresponds to $(1, 1)$. When a point is deviated to other quadrants, the chip error occurs. Define $d_{i,j}^{(k)}$ as the amplitude of EVM of kth chip point to the standard location (i, j) and $p_{i,j}^{(k)}(i = 0, 1, j = 0, 1, k = 0, 1, \ldots, l/2)$ as the probability of each decoding case, where

$$p_{i,j}^{(k)} = \frac{\frac{1}{d_{i,j}^{(k)}}}{\sum\limits_{r=0,1,t=0,1} \frac{1}{d_{r,t}^{(k)}}}. \tag{4}$$

We use P_s to represent the possibility of sending some symbol s. The final probability of sending symbol s is the multiplication of the probability of all the decoding cases:

$$P_s = \sum_{k=1}^{l/2} p_{i,j}^{(k)}. \tag{5}$$

For instance, the possibility of sending symbol 0100 is $P_4 = p_{0,1}^{(0)} p_{0,0}^{(1)}$. The possibility of sending other symbols can also be acquired similarly. Thus we get P_s for $s = 0, 1, \ldots, 15$. We normalize the P_s into $\mathcal{P}_s$, we call which *bit confidence*, to represent the possibility that transceiver sends a symbol s:

$$\mathcal{P}_s = \frac{P_s}{\sum_{j=1}^{n} P_j}. \tag{6}$$

Suppose the RX has decoded the chip sequence into some symbol s', the confidence of the coded symbol is exactly represented by $\mathcal{P}_{s'}$. Finally, for all the symbols in the packet, the overall BER can be calculated by

$$\text{BER}_u = \frac{\sum_{i=1}^{L}(1 - \mathcal{P}_{s'}^{(i)})}{L}, \tag{7}$$

where L is the packet size.

2) Calibration: Although the above likelihood-based estimation runs well in most cases, we observe that BER is over-estimated in low-quality links. One factor is the increase of 'ambiguous chips'. Ambiguous chips is a point disperses to location hard to distinguish between '0' and '1'. These chips have a severe impact on BER estimation because when the RX runs ML algorithm to find the most probable symbol, there may be non-ground-truth symbol taking non-negligible confidence. As a result, the confidence of ground-truth symbol is decreased even though the decoding causes no bit errors. To remove the effect of ambiguous chips, we multiply a factor α to obtain the final BER. In our design, we assume α a function of the spreading length l and SNR, thus we make extensive experiments to test which value performs best given the spreading length and SNR. Since SNR estimation suffers inaccuracy, we use CER instead as a parameter to obtain proper α. A traditional approach to get CER is counting the chip differences between the received chip sequence with the standard sequence after decoding, it is inadvisable here because a chip sequence can be incorrectly decoded.

Mrs.Z closes the gap by theoretically calculating the possibility of each sequence to be incorrectly decoded. Specifically, with two symbols s_1 and s_2, let $\text{D}(\cdot)$ be the DSSS function and r be the received chip sequence. Intuitively, if the hamming distance of two codewords $H(\cdot)$ is larger, the possibility of mis-decoding between these two symbols is lower. We use $P_h(s_i, s_j)$ to represent the probability of mis-decoding between s_i and s_j. Thus

$$P_h(s_i, s_j) = f(|H(\text{D}(s_i), \text{D}(s_j))|), \tag{8}$$

where $f(\cdot)$ is the probability density function (PDF) of Gaussian Distribution $\mathcal{N}(0, \sigma_1^2)$. After normalizing $P_h(s_i, s_j)$ to $\mathcal{P}_h(s_i, s_j)$, the estimated chip errors in decoding chip sequence r is given as:

$$N_{\text{chip\_error}} = \sum_{i=1}^{16} \mathcal{P}_h(s, s_i) \times H(r, \text{D}(s_i)), \tag{9}$$

where s is the symbol decoded from r. The above procedures ensure that even in a poor channel quality, CER estimation will not be severely affected by mis-decoding. In our experiments, we divide CER to four different levels and explore the most effective α for each length and each CER level with trial-and-error so that BER can be estimated with no severe deviation.

3) Cross-Length BER Estimation: Next, the RX makes cross-length BER estimation to compare which spreading length corresponds to the highest throughput. To estimate BER of other lengths except from the current one, we leverage SNR as a bridge. Specifically, the RX obtains the SNR corresponding to the BER in the current length so that it obtains the BER of other lengths based on SNR-BER relation. However, one problem is that BER-SNR relation suffers variations that will cause the RX making the wrong decision, as shown in Fig. 8.

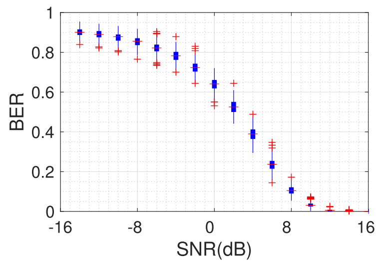

Fig. 8. BER-SNR variations when the spreading length is 8.

To avoid unintended packet loss, we take a conservative scheme to exploit maximum and minimum BER in each length as the ruler for length changing. We observe that BER of other lengths always ranges in certain intervals if we calculate from BER-SNR relation. Let the range be $(b_{\min,l}, b_{\max,l})$ for spreading length l. We assume that if the mean throughput of the spreading length $\frac{l}{2}$ is higher even with BER $= b_{\max,\frac{l}{2}}$, the spreading length should change from l to $\frac{l}{2}$. For the length change from l to $2l$, the rule is similar.

4) Rate Selection: After cross-length BER calculation, Mrs.Z obtains the BER fluctuation range $(b_{\min,l}, b_{\max,l})$ for the spreading length l so that the RX can make rate selection. The principles of rate selection can be summarized as:

- The data rate never increases unless the higher rate calculated with the maximum BER outperforms the current rate;
- The data rate drops when a lower rate calculated with the maximum BER outperforms the current rate.

To illustrate the effects of BER on throughput, we define the effective throughput T_{eff} in Mrs.Z:

$$T_{\text{eff}} = v(1 - b), \tag{10}$$

where v is theoretical data rate. In our rate adaptation scheme, the receiver sends a rate-increasing signal only if higher throughput can be obtained. With pre-computed BER of the current length L_i, Mrs.Z regulates the rate increasing condition as follows:

$$v_l(1 - b_l) < v_{\frac{l}{2}}(1 - b_{\max,\frac{l}{2}}). \tag{11}$$

Similarly, for rate lowering, the condition is regularized as:

$$v_l(1 - b_l) < v_{2l}(1 - b_{\max,2l}). \tag{12}$$

Mrs.Z allows jumping over multiple rates from a higher to a lower rate, e.g., from 4-to-4 to 4-to-64 spreading. The inverse case, however, is prohibited. Switching to a higher rate is more gradual to avoid retransmission caused by accidental inaccurate estimation.

C. Enhancement

For now, Mrs.Z has built the framework for efficient multi-rate transmission. However, several problems in real-link transmission remain to be addressed. The first is inefficient retransmission. Traditionally, the TX has to retransmitted the whole packet repeatedly when the CRC validation in the RX fails. This scheme restricts the performance of Mrs.Z severely because the TX sometimes should frequently select the larger L_s conservatively to bypass retransmission cost. Another problem is the inaccuracy estimation of BER caused by data outlier. Data outlier occurs in many cases, e.g., instant distortion in electromagnetic signal and moving objects. Some outlier persists for less than one-packet transmission time, but they lead to the inaccuracy in BER estimation for the next transmission.

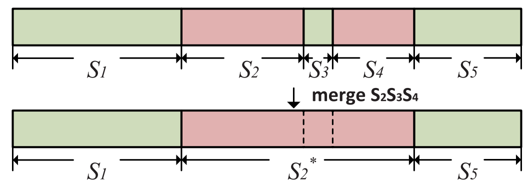

Fig. 9. Merging sacrifices retransmission cost to reduce encoding cost. Mrs.Z quantizes them to achieve efficient encoding and retransmission.

In this section, we add two schemes, partial retransmission and outlier removal, in the framework to maximize Mrs.Z's performance.

1) Partial Retransmission: We inherit the smart retransmission [28] to bypass overly conservative rate selection. Smart retransmission efficiently supports partial retransmission by encapsulating the index of low-confidence bits in the feedback link. Thus the TX only needs to retransmit these bits instead of the whole packet.

Without loss of generosity, a packet with length L takes $\log L$ bits to encode the index of a bit. When low-confidence bits are dispersive, merging consecutive low-confidence bit indexes is more efficient as shown in Fig. 9. To encode efficiently, Mrs.Z quantizes the encoding cost and retransmission cost in dynamic programming and judge whether a merging is required [29]. Mathematically, when a data packet comes, the receiver decodes it to a sequence of symbols $\{s_i\}$ and correctly-despreading confidence $\{c_i\}$. With a trained threshold Th, the bits are categorized into low-confidence and high-confidence symbols. Then it forms a sequence as follows:

$$S_1^h S_1^l S_2^h S_2^l \cdots S_k^h,$$

where S_i^h is the ith slices in which all bits have high confidence and S_i^l is the one in which all bits have low confidence. It begins and ends with a high-confidence slice whose length can be 0. If the cost of encoding S_{i-1}^l and S_i^l exceeds the retransmission cost of S_i^h, the RX merges the three slices $(S_{i-1}^l, S_i^h, S_i^l)$. Let the maximum packet length be L, without merging, the cost is:

$$C^n(i-1, i) = 3\log(L) + (S_{i-1}^l + S_i^l) + \sigma_n, \tag{13}$$

where the first term is encoding cost, the second term is retransmission cost and σ_n is a constant. With merging, the cost is:

$$C^m(i-1, i) = \log(L) + (S_{i-1}^l + S_i^h + S_i^l) + \sigma_m. \tag{14}$$

When $C^m(i-1, i) - C^n(i-1, i) < 0$, a merging operation is required. Thus, the true cost is

$$C(i-1, i) = \min\{C^n(i-1, i), C^m(i-1, i)\}. \tag{15}$$

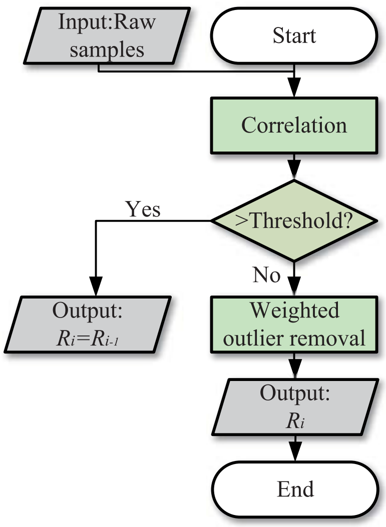

Fig. 10. The flow diagram of correlation-and-weighted outlier removal.

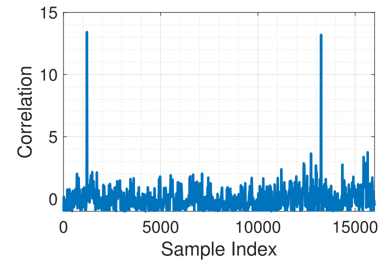

Fig. 11. The normalized correlation spikes when the preamble is aligned. If multiple spikes are detected, the RX assumes a collision occurs.

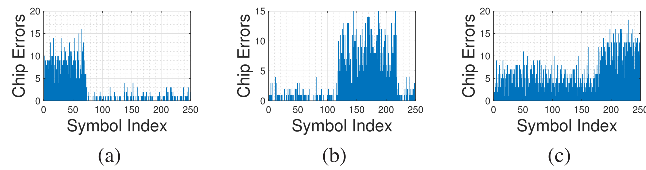

Fig. 12. Observable physical hops can occur in the head, body, or tail and their persisting time is variable. (a) Packet-head hops. (b) Packet-body hops. (c) Packet-tail hops.

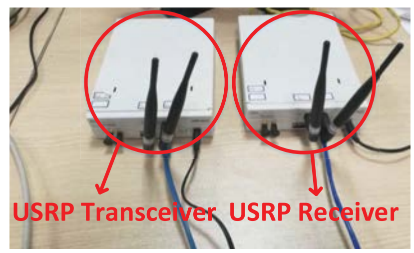

Fig. 13. Our testbed uses two USRPs as the ZigBee transceiver and receiver.

2) Outlier Removal: Although the framework of Mrs.Z for now runs well in most cases, Mrs.Z suffers severe degradation in rate adaptation when outliers occur. Deviating dramatically from the normal, these outliers have significant effect on calculating BER. However, they should not be considered for rate changing since these outliers do not reflect the link conditions.

In the complicated industrial wireless links, we observe there are two main causes for outlier:

- **Collisions**. The collision produced by concurrent transmission from multiple TXs dramatically increases the chip errors in overlapping segments. Without a special decoding scheme, the RX simply assumes the overlapping segments as errors and attributes the spike of errors to poor link conditions. As a result, the RX will return an incorrect feedback to the TX to lower the rate.
- **Physical hops**. Physical factors such as moving objects, electromagnetic interference and hardware imperfection have non-negligible effects on BER. Since such effects are mostly temporal, we can observe a 'hop' (sudden increase or decrease) in the chip errors. Physical hops lead to the dramatic changes in chip errors and BER in certain time. However, the duration of such hops is unpredictable, where it can persist across several consecutive packets, or terminates in one packet. Neglecting the effect of hops leads to the waste of the bandwidth since hops significantly increase the BER as a probe to lower the rate, even if some of them terminate quickly and have no effect on the next packet.

To detect and distinguish collisions and physical hops, we propose a two-step *correlation-and-weighted outlier removal* (COR) as an enhancement for Mrs.Z. The block diagram of COR is shown in Fig. 10. First, the RX exploits the known pattern in the preamble to make cross-correlation to detect if there is a collision. The details of cross-correlation can be referred in [7]. Mathematically, the correlation spikes if the correlation starts exactly at the first sample of the preamble. If multiple spikes are detected in the packet, as shown in Fig. 11, the RX assumes a collision occurs and notifies the TX to retain the original data rate.

Second, if there is no collision, the RX further seeks for physical hops in the received packet in terms of the chip errors in each symbol. According to the hop location, we roughly classify the physical hops into packet-head hops, packet-body hops and packet-tail hops, as shown in Fig. 12. For packet-head and packet-body hops, the sudden error increasing terminates before the end of the packet, which means these physical hops have little chance to influence the next packet. COR assigns a lower weight to the errors in such hops to reduce their impact on BER calculation. According to the simulation, a weight of 0.1$\sim$0.3 performs well. For packet-tail hops, the outlier manifests itself as a sudden increase with no decrease until the end of the packet and there is much possibility the hop will persist to the next packet. To ensure the reliability of next transmission, COR assigns a higher weight to the errors in the packet-tail hop, which should be > 0.7 empirically.

V. Implementation

As shown in Fig. 13, we have implemented Mrs.Z on two USRP N210s based on the 802.15.4 framework from [30].

Physical Layer: Mrs.Z makes lightweight modification on the traditional ZigBee framework at the transceiver. The output of bit-to-symbol module is sent to the DSSS selector, where Mrs.Z reads the advertised spreading length in the last

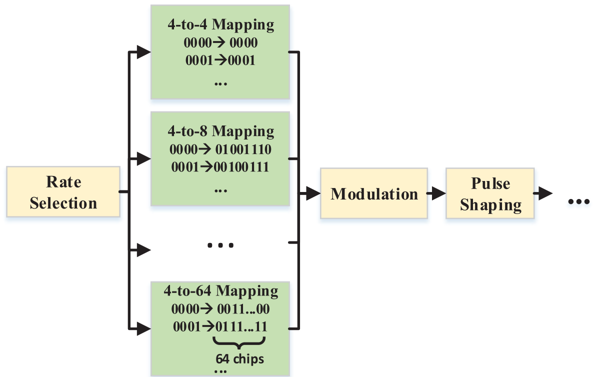

Fig. 14. Rate selection in the transceiver.

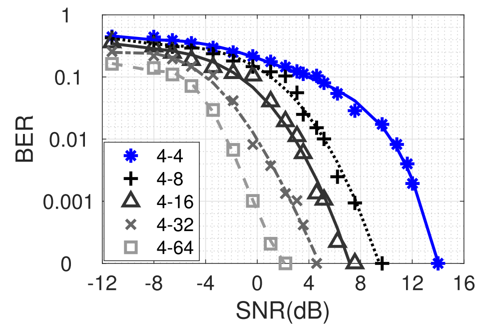

Fig. 15. The general BER-SNR relation when symbols are spread to different lengths.

feedback received from the receiver, as shown in Fig. 14. If no feedback can be read, the transceiver uses the 4-to-64 spreading length, the one with the lowest rate but the highest reliability, in this transmission. After transmission, the transceiver enters a state of waiting for feedback. If a feedback is received within a given time, the transceiver updates the last feedback in the cache and generates a new data packet with the up-to-date advertised rate for transmission. Otherwise, it decreases its data rate and retransmits the last packet.

At the receiver, Mrs.Z keeps detecting the preamble. Once a packet is detected, Mrs.Z estimates the CER and BER respectively and further calibrates bit errors for rate selection. After removing the outliers, the receiver divides and merges the low-confidence segments and issues a feedback for smart retranmission.

Interference Emulator: Fine-grained interference control in real links is challenging. Instead, we emulate interferences by mixing the source signal with a Gaussian noise source in the GNURadio. We adjust the amplitude of noise from -12 to 14 dB to validate Mrs.Z's robustness to noisy environments.

Implementation Complexity: Mrs.Z customizes both the TX and RX with lightweight modification on the legacy ZigBee devices to support multi-rate transmission. In the TX, what Mrs.Z needs is simply selecting the proper rate from the cache and encoding the bit stream into the pre-defined chip sequence. In the RX, the decoding and the rate-selection work concurrently. The decoding corresponds to the modification in the TX, where Mrs.Z decodes the data rate in the spreading length field and selects the proper chip-symbol mapping. For the rate selection, we implement CER estimation, BER estimation, BER calibration and rate selection in a different thread to boost Mrs.Z's parallelism. Overall, Mrs.Z has the advantage of modularity, lightweight modification and parallelism, which makes it feasible in implementation.

VI. Experiments

Based on our testbed, we conduct extensive experiments to evaluate the performance of Mrs.Z in a general office scenario.

A. BER-SNR

We first plot the general BER-SNR relation to demonstrate that enabling multi-rate transmission in ZigBee results in different levels of resistance to noise when using different spreading lengths. We vary the SNR from -12dB to 14dB. Results have shown in Fig.15. Note that when SNR is smaller than -12dB, the receiver can hardly detect the preamble of the packet. Results show that $l = 64$ is the robustest that can keep the BER under 3×10^{-3} with -2dB SNR. When l decreases from 64, while the TX sends at a higher data rate, the capability of error-avoidance is degrading.

B. CER Measurement Accuracy

We evaluate Mrs.Z on the accuracy of CER measurements. The accuracy of CER determines whether BER estimation can be correctly calibrated. To make evaluation, we let the transceiver transmit data every 500 milliseconds in both static and mobile channels. The channel mobility is made by (i) moving the transceiver at a slow speed (0-0.2m/s); (ii) moving the transceiver at a high speed (0.6-0.8m/s); (iii) inserting a moving iron board between the transceiver and the receiver. We measure Mrs.Z's performance across different transmission power, transceiver positions, SNRs (i.e, by mixing a noise source to the signal source), and spreading lengths. For each transmission, the receiver computes CER with chip errors divided by chip number, where chip errors are obtained from Eq. 9.

We compute the ground-truth CER by comparing the received sequence with the transmitted data. Fig. 16a compares the estimated CERs with their ground truth in static channels. If chip errors can hardly be observed in a single packet, chip errors in several consecutive packets are aggregated to estimate overall CER.

To show the necessity of Mrs.Z, we conduct two contrast methods for CER measurements.

- **Estimating with Chip Differences:** The receiver will decodes the received chip sequence to a symbol directly without considering whether this symbol is the correct one. Then, chip differences between the decoded symbol and the received sequence are assumed as chip errors directly. Apparently the approach is zero-error in high SNRs but degrades drastically as the SNR decreases. The result is shown in Fig. 16b.
- **Estimating with Preamble Probe:** The receiver estimates chip errors with the error counts in the preamble. Because the preamble is fixed and uses 4-to-64 spreading, the chip error estimation in the preamble is almost

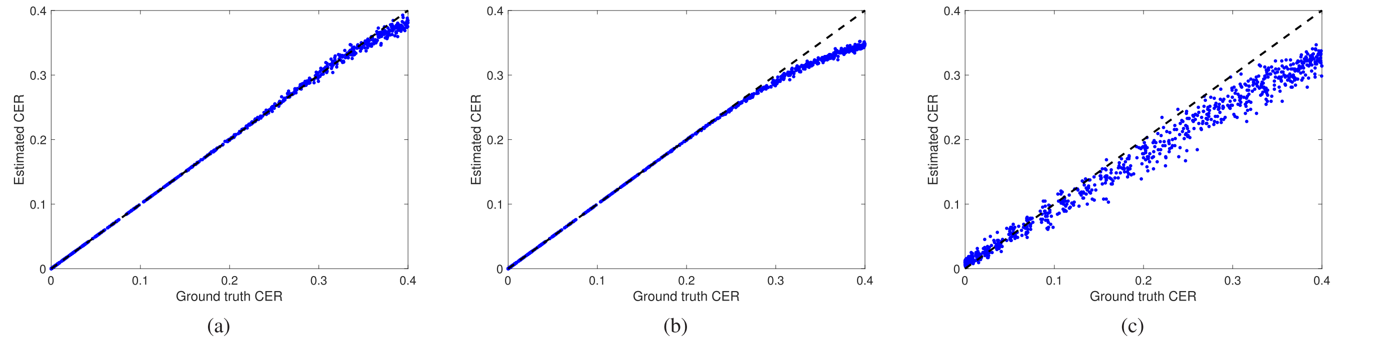

Fig. 16. Evaluation of CER estimation. (a) CER estimation with Mrs.Z. In low SNR parts, Mrs.Z can reduce the impact of under-estimation. (b) CER estimation with chip differences. (c) CER estimation with preambles.

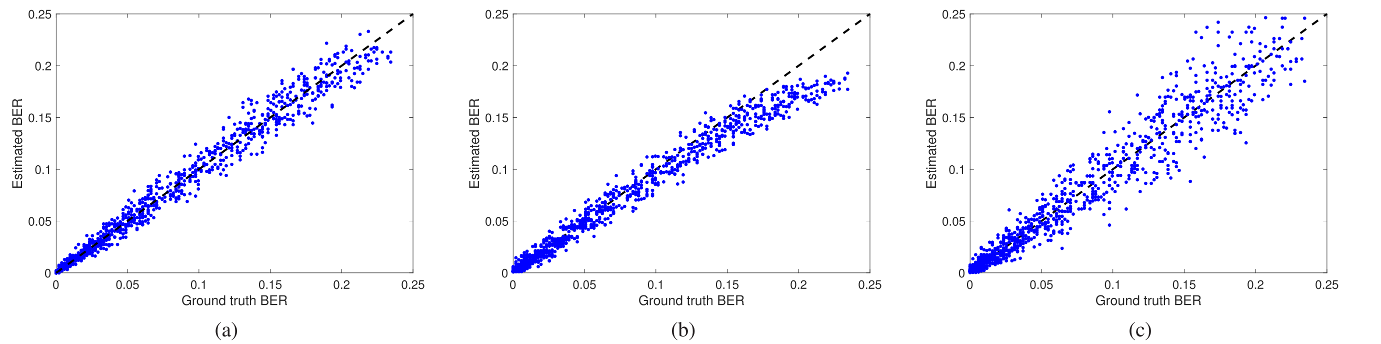

Fig. 17. Evaluation of BER estimation. (a) BER estimation with Mrs.Z. (b) BER estimation with SoftRate. (c) BER estimation with preambles.

zero-error. However, as known that the preamble of ZigBee is rather short (4 bytes), thus it will lead to the severe dispersion when it is used to estimate the chip errors in the whole packet. The result is shown in Fig. 16c.

Results show that Mrs.Z (Fig. 16a) can have a highly accurate estimation on CER compared to benchmarks (Fig. 16b, 16c). When CER reaches 0.4 or higher, the under-estimation occurs even with Mrs.Z. However, it is acceptable because when CER> 0.4, the RX will always select 64 chips as the spreading length that the RX does not need CER to calibrate the deviation in BER. Compared to Mrs.Z, both the chip-difference and preamble-based estimation suffer severer underestimation and the preamble-based estimation has large dispersion across all SNRs.

C. BER Estimation Accuracy

We test Mrs.Z's performance in predicting BER in both static and mobile channels as stated in Section VI-B. The transceiver generates and transmits packet every 500 milliseconds. A simulated noise source is added to the original signal. SNR ranges from -8dB to 14dB. In the both static and mobile channels, the transceiver transmits about 1000 packets. The payload length is fixed to 64 bytes. The RX maps raw samples into constellation and estimate each chip's confidence. Then the BER is estimated as stated in Section IV-B. To obtain the ground truth BER, the receiver compares the despread symbols with original symbols which we store at the receiver. We show the result in Fig.17a. To make a contrast, we use two benchmarks for BER estimation.

- **SoftRate:** SoftRate [12] has provided a method of exploiting LLRs to estimate overall chip errors. With no explicit approaches to mapping overall chip errors to bit errors, we build an empirical mapping between BER and CER. Note that we cannot calculate each chip's confidence as we do in Mrs.Z because SoftRate only gives an overall estimation on CER. The result is shown in Fig. 17b.
- **Estimation with Preamble Probe:** The receiver exploits the preamble field in the header as an SNR estimator. A generally-taken approach is calculating the error dispersion in the preamble to estimate SNR and further obtain BER. In addition to the short preamble, this approach also suffers unstable SNR-BER relation. The result is shown in Fig. 17c.

The results plotted in Fig. 17a to Fig. 17c show that Mrs.Z outperforms the other two approaches both in the static and mobile channels. For the first approach exploiting CER value for BER estimation, the errors come from (i) non-constant mapping factor from CER to BER, and (ii) SoftRate is a codeword-free approach neglecting the characteristics of DSSS. In contrast, Mrs.Z traverses the 16 standard chip sequences and calculate their confidence chip-by-chip to obtain the bit errors. For the SNR-based approach, we can see the variance effect of uncertain SNR-BER relation. Although dispersion in the preamble is highly accurate, using BER in the

TABLE I
LIST OF THE THROUGHPUT IN DIFFERENT SNRS

Spreading length	>13dB	9-13dB	5-9dB	0-5dB	-4-0dB
4-to-4	157.2kb/s	68.7kb/s	16.5kb/s	0	0
4-to-8	106.3kb/s	105.1kb/s	54.9kb/s	12.2kb/s	0
4-to-16	69.2kb/s	67.9kb/s	64.0kb/s	18.8kb/s	4.8kb/s
4-to-32	36.4kb/s	35.2kb/s	34.8kb/s	34.3kb/s	11.7kb/s
4-to-64	19.4kb/s	19.7kb/s	19.2kb/s	19.3kb/s	19.3kb/s

header as a hint results in an even-worse performance. In our experiments, Mrs.Z is not responsive to the channel mobility in BER estimation. Although the overall BER increases, the accuracy of estimation stays stable.

D. Throughput

In this section, we evaluate the throughput performance of Mrs.Z. We first look for the optimal transmission rate in different SNR values. We divide several intervals of SNR and make the transceiver always transfer at a fixed spreading length. With the average data rate and BER, we can get the effective data rate. We further demonstrate Mrs.Z's effectiveness by comparing the throughput of Mrs.Z against the traditional ZigBee and SoftRate. The experiments in this section are conducted in a laboratory environment, thus we emulate noise and interferences using a noise source in GNURadio. Again, we measure their performance in both static and mobile channels. Mobile channels are emulated as stated in Section VI-B.

Throughput in Different SNRs: We evaluate the throughputs with different spreading lengths in different SNRs. We divide SNR into five intervals and make evaluation for each spreading length respectively. Results in Table I show that each spreading length outperforms others in the certain SNR interval. When SNR is smaller than -4dB, the throughput is extremely low even if using 4-to-64 spreading. Another fact we observe is that even the spreading length is truncated half, the effective rate can hardly obtain a 2x performance in throughput. Considering the noise and interference, using a high rate means decreases the self-defending capability to noise and further causing an increase in bit errors. In the cases of high-quality channels or large spreading lengths with strong defending capability, the optimal coefficient can almost achieve 2 ($\frac{36.4}{19.4} = 1.87$ compared to $\frac{157.2}{106.3} = 1.48$).

Methodology and Benchmarks: We first present the methodology of performance comparison. The send-feedback mechanism is as follows: once the transceiver transmits a packet, it enters the waiting state. The receiver feedbacks a short packet to the transceiver after receiving a packet.

When the transceiver is in a waiting state, it changes its role to a receiver waiting for the reception of a reply. If a feedback with an ACK is received, it triggers a new transmission. On the other hand, if a feedback with an NAK is received or the waiting time exceed, it triggers a retransmission. The timeout threshold here is set to 4x of the time to transmit a standard full-payload ZigBee packet. The payload length is set to 64 bytes in the experiments. To verify Mrs.z's throughput, we emulate channel with varying SNR in a laboratory environment.

We choose traditional ZigBee, SoftRate, and Smart-Pilot [14] as benchmarks for throughput evaluation.

- **Traditional ZigBee:** The transceiver spreads each symbol to a 32-chip sequence. The theoretical data rate is fixed at 250kbps.
- **SoftRate:** To apply SoftRate, we firstly implement multi-rate transmission as designed in Mrs.Z, and then we use SoftRate for rate adaptation. SoftRate adds a PHY-MAC interface called SoftPHY for BER estimation. It exploits LLRs [12] as physical hints to obtain estimated BER across the packet. SoftRate is designed originally in WiFi but can be transferred to ZigBee with little modification. Since a ZigBee RX receives chips before bits, SoftRate estimates CER first and then uses CER to estimate BER. Compared to the legacy ZigBee, SoftRate's major overhead comes from the design of SoftPHY, which provides an extra interface for the upper call.
- **Smart-Pilot:** Similarly, we support multi-rate transmission first and implement Smart-Pilot. Smart-Pilot combines decoded bits with relatively high confidence level and 802.15.4 protocol headers as the pilots to calibrate CSI. Using LLRs to obtain rough BER, Smart-Pilot exploits pilots to remove the residual channel effect. The threshold of determining a high-confidence bit is set to 0.8 empirically. In ZigBee, the length of the preamble is 4 bytes. Smart-Pilot's overhead mainly lies in computing of high- and low-confidence bits in order respectively.
- **Mrs.Z:** The receiver computes and calibrates BER with the method provided in Section IV and uses partial retransmission and outlier removal for further improvement. The receiver would issue a feedback packet recording the low-confidence part. Note that this only works when the feedback packet is lightweight. For each data packet, the receiver also attempts to detect whether there are outliers that would have significant effect on rate selection. The outlier removal procedures are as mentioned in Section IV-C. As for the complexity, Mrs.Z is a bit more complicated than SoftRate and Smart-Pilot with the outlier removal schemes. However, the purpose of Mrs.Z is to propose a multi-rate selection scheme robust in all kinds of scenarios, which focuses more on reliability in the trade-off with simplicity. Therefore, reasonable increase in complexity is tolerable.

Performance: In emulated channels, a variable Gaussian-noise source is added to the signal source. We adjust the mean value and variance of the Gaussian noise to observe the ZigBee's adaptability when applied different rate selection schemes. SNR is set to step from -8dB to 14dB. In slow-mobility scenarios, transceivers are moved in a low and stable speed and in fast-mobility scenarios the transceivers are moved more drastically. The mean throughput in each rate selection scheme is shown in Fig. 18.

From the figure, Mrs.Z outperforms other three designs in static and slow-mobility scenarios. Mrs.Z outperforms others because both SoftRate and SmartPilot are weak in BER estimation if transplanted into ZigBee directly. SoftRate has to rely on the fixed CER-BER transition to obtain BER. However, due to the difference in chip error distribution, BER can be

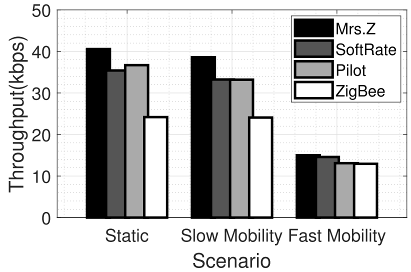

Fig. 18. Mean throughput of different schemes.

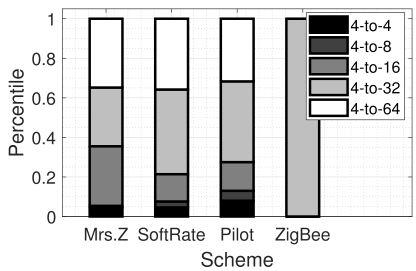

Fig. 19. Distribution of selected spreading lengths in static channel.

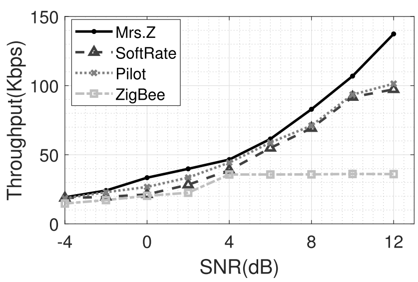

Fig. 20. The impact of SNR on throughput in the static channel.

quite different even with the similar CER. SmartPilot attempts to use the soft and hard pilot to calibrate BER, but such pilots are scarce in ZigBee with the finite preamble and payload.

The throughputs achieved with all the four methods decrease significantly in the mobile channel. In fast-mobility scenarios, silent packet loss occurs frequently, where the receiver cannot even detect the preamble field. Even a packet is detected, it is always the case that 64 chips are used with whichever method.

To further study the performance variance, we track the ratios of all spreading lengths selected by each scheme in the static scenario, as shown in Fig. 19. For traditional ZigBee, the improvements lie in two aspects. On the one hand, Mrs.Z can select a higher rate in a good-quality channel. On the other hand, the fixed-spreading-length of traditional ZigBee is more likely to cause retransmission when the noise is large, while Mrs.Z can select a lower rate to avoid retransmission, improving the throughput. Compared to SoftRate and SmartPilot, Mrs.Z selects $l = 16$ more often. The differences in rate selection are also caused by the inaccuracy of SoftRate and SmartPilot in BER prediction in ZigBee.

We also track the variations of mean throughput when SNR varies, as shown in Fig. 20. We adjust the parameters of the noise source to set SNR from -4dB to 12dB. Results show that Mrs.Z outperforms other methods across the SNR interval when applied in ZigBee.

TABLE II

PART OF RESULTS ON THE IMPACT OF SNR AND THE LOCATION OF THE HIDDEN TERMINAL

Distance	0.5m	1m	2m	3m	4m
0dB	85.03%	98.21%	77.16%	88.49%	93.85%
-4dB	67.06%	97.71%	69.61%	80.23%	84.30%
-8dB	61.53%	86.01%	63.55%	72.67%	78.96%
-12dB	41.23%	41.07%	48.94%	45.12%	43.89%

E. Outlier Removal

In this section, we evaluate the performance of COR in outlier removal. Mrs.Z's outlier removal strategy mainly avoids unintended rate decrease resultant from i) collisions and (ii) unpredicted physical hops. In the experiment, we have two noise sources N_1, N_2, one is used to generate constant noise (N_1) and the other is used to generate physical hops (N_2). We make fine-grained adjustments on the properties of N_2, including amplitudes, persisting time, etc., to observe to what extent the enhancements would improve the throughput performance. We set the distance between TX and RX to $1m$. We also use another USRP N210 as a hidden terminal that generates packets regularly to make collisions. If the chip errors increase sharply, this symbol will be marked as the head of a candidate outlier segment. Aligning with the first sample, the receiver calculates correlation and judges whether the increasing errors is from a multi-packet collision by comparing the correlation with a threshold. If it is a collision, the receiver dumps the packet and makes no change on the data rate. Otherwise, the receiver removes the impact of outliers as illustrated in Section IV-C.

We firstly evaluate the accuracy in rate selection when COR is implemented. The outliers are tightly related to the noise source and the hidden terminal. Hence, we assume SNR and the distance of interference packets as potential parameters and change them to observe their impact on the accuracy. When the receiver decodes a packet, it applies COR and determines the data rate of the next packet. When the next packet arrives, the receiver can examine whether it is the proper data rate by comparing the current effective throughput and the theoretical effective throughput of other lengths. If the current effective throughput outperforms all the others, we assume it a *hit*, otherwise we assume as a *loss*. We quantify the accuracy with

$$\beta = \frac{n_h}{n_h + n_l}, \tag{16}$$

where n_h, n_l are the total number of hits and losses, respectively.

The results are listed in Table II. From the results, COR performs better when SNR increases. We suppose the observation is mainly resultant from two reasons. i) Lower SNR makes the correlation fluctuate dramatically and leads to the mis-determination. The index with spike in a collision is more often taken as a non-spike index, or reversely. ii) As SNR increases, the packet is more likely to mis-select the data rate, no matter whether the COR is applied. The effect of the distance of the hidden terminal is a bit more complicated. Results show that when the distance between the hidden terminal and the RX, d_{H-RX}, is equal to or much larger than

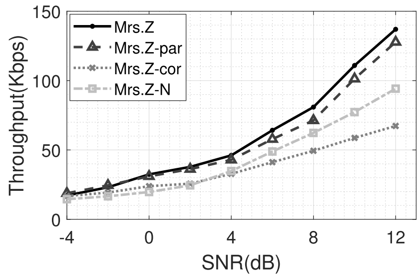

Fig. 21. The impact of our enhancements on throughput.

the distance between the TX and the RX, d_{TX-RX}, COR can achieve its best performance. When $d_{I-RX} = d_{TX-RX}$, the correlation does not change much, and the spike stands out remarkably. Thus it is easy for the RX to distinguish whether a collision occurs. When $d_{I-RX} \gg d_{TX-RX}$, the interference can be simply treated as low-amplitude noise, which has little effect on the rate selection.

We further study the impact of outlier removal on throughput. In the experiment, the properties of N_2 and the hidden terminal change randomly and N_1 is controlled by adjust the value of SNR. The results are shown in Fig. 21, where we use Mrs.Z, Mrs.Z-par, Mrs.Z-cor, Mrs.Z-N to represent the case of implementing both partial retransmission and COR, only partial retransmission, only COR and neither of partial retransmission nor COR respectively. From the results, our enhancements achieve significant effects, especially in high-quality links. Compared with Mrs.Z with no enhancements, the results of implementing partial retransmission and outlier removal achieve an improvement of 87.8%. We can draw another conclusion from the figure that, in normal channels, partial retransmission is more influential than outlier removal. With partial retransmission, the transceiver can avoid over-conservative selection of data rates. Outlier removal is less impactive because of its low occurrence rate and randomness.

VII. Conclusion

In this paper, we have proposed Mrs.Z, a physical design enabling multi-rate transmission and rate selection in ZigBee. We first leverage the inherent DSSS module in ZigBee to make ZigBee able to select different rates for transmission, which only needs lightweight modifications. Then, a rate selection scheme feasible on ZigBee is proposed in Mrs.Z. By leveraging the characteristics of DSSS, Mrs.Z makes fine-grained estimation on BER and selects the proper rate by calculating effective throughput. We implement Mrs.Z and verify its effectiveness on USRPs. Results show that Mrs.Z achieves an improvement of 15%, 20% and 80% over Smart-Pilot, SoftRate and traditional ZigBee in throughput on average.

References

[1] Y. Cao *et al.*, "Mrs. Z: Improving ZigBee throughput via multi-rate transmission," in *Proc. ICNP*, Oct. 2017, pp. 1–10.

[2] (2018). *ZigBee Standard*. [Online]. Available: http://standards.ieee.org/getieee802/download/802.15.4-2011.pdf/

[3] H. Y. Tung *et al.*, "The generic design of a high-traffic advanced metering infrastructure using ZigBee," *IEEE Trans. Ind. Informat.*, vol. 10, no. 1, pp. 836–844, Feb. 2014.

[4] P. Yi, A. Iwayemi, and C. Zhou, "Developing ZigBee deployment guideline under WiFi interference for smart grid applications," *IEEE Trans. Smart Grid*, vol. 2, no. 1, pp. 110–120, Mar. 2011.

[5] Y. C. Tay, K. Jamieson, and H. Balakrishnan, "Collision-minimizing CSMA and its applications to wireless sensor networks," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 6, pp. 1048–1057, Aug. 2004.

[6] M. Doddavenkatappa, M. C. Chan, and B. Leong, "Splash: Fast data dissemination with constructive interference in wireless sensor networks," in *Proc. USENIX NSDI*, 2013, pp. 269–282.

[7] S. Gollakota and D. Katabi, "ZigZag decoding: Combating hidden terminals in wireless networks," in *Proc. ACM SIGCOMM*, 2008, pp. 159–170.

[8] L. Kong and X. Liu, "mZig: Enabling multi-packet reception in ZigBee," in *Proc. 21st Annu. Int. Conf. Mobile Comput. Netw.*, Sep. 2015, pp. 552–565.

[9] (2018). *802.11n Standard*. [Online]. Available: http://standards.ieee.org/getieee802/download/802.11n-2009.pdf/

[10] J. C. Bicket, "Bit-rate selection in wireless networks," Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci., MIT, Cambridge, MA, USA, 2005.

[11] S. H. Y. Wong, H. Yang, and S. Lu, "Robust rate adaptation for 802.11 wireless networks," in *Proc. 12th Annu. Int. Conf. Mobile Comput. Netw.*, Sep. 2006, pp. 146–157.

[12] M. Vutukuru, H. Balakrishnan, and K. Jamieson, "Cross-layer wireless bit rate adaptation," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 4, pp. 3–14, Aug. 2009.

[13] G. Judd, X. Wang, and P. Steenkiste, "Efficient channel-aware rate adaptation in dynamic environments," in *Proc. ACM MobiSys*, Jul. 2008, pp. 118–131.

[14] L. Wang *et al.*, "Wireless rate adaptation via smart pilot," in *Proc. ICNP*, Oct. 2014, pp. 409–420.

[15] P. A. K. Acharya *et al.*, "Congestion-aware rate adaptation in wireless networks: A measurement-driven approach," in *Proc. 5th Annu. IEEE Commun. Soc. Conf. Sensor, Mesh Ad Hoc Commun. Netw.*, Jun. 2008, pp. 1–9.

[16] A. Kamerman and L. Monteban, "WaveLAN-II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Tech. J.*, vol. 2, no. 3, pp. 118–133, May 1997.

[17] J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs," in *Proc. INFOCOM*, Apr. 2006, pp. 1–9.

[18] J. Zhang, K. Tan, J. Zhao, H. Wu, and Y. Zhang, "A practical SNR-guided rate adaptation," in *Proc. INFOCOM*, Apr. 2008, pp. 2083–2091.

[19] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 packet delivery from wireless channel measurements," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 22, no. 4, pp. 159–170, 2010.

[20] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi, "AccuRate: Constellation based rate estimation in wireless networks," in *Proc. 7th USENIX Conf. Netw. Syst. Design Implement.*, Apr. 2010, pp. 175–190.

[21] K. D. Huang, K. R. Duffy, and D. Malone, "H-RCA: 802.11 collision-aware rate control," *IEEE/ACM Trans. Netw.*, vol. 21, no. 4, pp. 1021–1034, Aug. 2013.

[22] J. Huang, Y. Wang, and G. Xing, "LEAD: Leveraging protocol signatures for improving wireless link performance," in *Proc. 11th Annu. Int. Conf. Mobile Syst., Appl., Services*, Jun. 2013, pp. 333–346.

[23] G. Wang, S. Zhang, and K. Wu, "TiM: Fine-grained rate adaptation in wlans," in *Proc. ICDCS*, Jul. 2014, pp. 577–586.

[24] D. Giustiniano, D. Malone, D. J. Leith, and K. Papagiannaki, "Measuring transmission opportunities in 802.11 links," *IEEE/ACM Trans. Netw.*, vol. 18, no. 5, pp. 1516–1529, Oct. 2010.

[25] I. Pefkianakis *et al.*, "MIMO rate adaptation in 802.11n wireless networks," in *Proc. 16th Annu. Int. Conf. Mobile Comput. Netw.*, Sep. 2010, pp. 257–268.

[26] W. L. Shen, K. C. J. Lin, and S. Gollakota, "Rate adaptation for 802.11 multiuser MIMO networks," *IEEE Trans. Mobile Comput.*, vol. 13, no. 1, pp. 35–47, Jan. 2014.

[27] C. Y. Li, C. Peng, and P. Cheng, "An energy efficiency perspective on rate adaptation for 802.11n NIC," in *Proc. ACM MOBICOM*, Aug. 2016, pp. 1–9.

[28] M. O. Khan, L. Qiu, and A. Bhartia, "Smart retransmission and rate adaptation in wifi," in *Proc. IEEE ICNP*, Aug. 2015, pp. 1–6.

[29] H. J. K. Balakrishnan, "PPR: Partial packet recovery for wireless networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 409–420, Aug. 2007.

[30] (2018). *ZigBee Implementation on GNURAdio*. [Online]. Available: https://github.com/bastibl/gr-ieee802-15-4/

Linghe Kong (S'09–M'13–SM'18) received the B.Eng. degree in automation from Xidian University in 2005, the master's degree in telecommunication from Telecom SudParis in 2007, and the Ph.D. degree in computer science from Shanghai Jiao Tong University in 2012. He is currently a Research Professor with the Department of Computer Science and Engineering, Shanghai Jiao Tong University. Before that, he was a Post-Doctoral Researcher with Columbia University, McGill University, and the Singapore University of Technology and Design. His research interests include wireless networks, big data, mobile computing, and the Internet of things.

Yifeng Cao received the bachelor's degree from Shanghai Jiao Tong University, in 2017. He is currently pursuing the M.S. degree with the Department of Computer Science and Engineering, Shanghai Jiao Tong University.

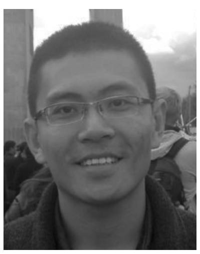

Liang He (S'09–M'12–SM'17) received the B.Eng. degree Tianjin University, China, in 2006, and the Ph.D. degree from Nankai University, China, in 2011. He is currently an Assistant Professor with the University of Colorado Denver (UCD). His research interests mainly focus on cyber-physical systems, IoTs, and mobile computing. Before joining UCD, he was a Research Fellow of The University of Michigan at Ann Arbor, Ann Arbor, MI, USA, a Research Scientist with the Singapore University of Technology and Design, Singapore, and a Research Assistant with the University of Victoria, Canada. He was a recipient of the Best Paper Award from the IEEE WCSP 2011, the IEEE GLOBECOM 2011, and QShine 2015, and the Best Paper Candidate of the IEEE GLOBECOM 2014.

Guihai Chen received the B.S. degree from Nanjing University in 1984, the M.E. degree from Southeast University in 1987, and the Ph.D. degree from The University of Hong Kong in 1997. He is currently a Distinguished Professor with Shanghai Jiao Tong University, China. He was a Visiting Professor with many universities including the Kyushu Institute of Technology, Japan, in 1998, the University of Queensland, Australia, in 2000, and Wayne State University, USA, from 2001 to 2003. He has a wide range of research interests with a focus on sensor network, peer-to-peer computing, and high-performance computer architecture and combinatorics.

Min-You Wu (S'84–M'85–SM'96) is currently a Professor with the Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China. He currently serves as the Chief Scientist of the Grid Center, Shanghai Jiao Tong University. He is a Research Professor with the University of New Mexico, USA. His research interests include wireless and sensor networks, vehicular networks, multimedia networking, parallel and distributed systems, and compilers for parallel computers. He has published over 200 journal and conference papers in his research areas. His research has been supported by the National Science Foundation, DoD, DoE, DARPA, the China 863 Program, the China 973 Program, Ministry of Education of China, and the Natural Science Foundation of China. He is a Senior Member of the CCF.

Tian He is currently a Full Professor with the Department of Computer Science and Engineering, University of Minnesota, Twin City. He received the Ph.D. degree under the supervision of Prof. J. A. Stankovic from the University of Virginia, VA, USA in 2004. He has the authored or coauthored over 200 papers in premier network journals and conferences with over 16 000 citations (H-Index 52). He was a recipient of the NSF CAREER Award, the George W. Taylor Distinguished Research Award, the McKnight Land-Grant Professorship, and five best paper awards in networking. He has served as a few program chair positions in international conferences and on many program committees. He has served as an Editorial Board Member of six international journals, including the *ACM Transactions on Sensor Networks* and the IEEE TRANSACTIONS ON COMPUTERS. His research interests include wireless sensor networks, cyber-physical systems, intelligent transportation systems, real-time embedded systems, and distributed systems, supported by the National Science Foundation, IBM, Microsoft, and other agencies.