Dense Deployment of BLE-based Body Area Networks: A Coexistence Study

Quang Duy La, *Member, IEEE,* Duong Nguyen-Nam, Mao V. Ngo, *Student Member, IEEE,* Hieu T. Hoang, and Tony Q.S. Quek, *Fellow, IEEE*

***Abstract*—The issues of cross-technology interference and coexistence in the unlicensed 2.4-GHz spectrum band among various technologies including WiFi, ZigBee, and classic Bluetooth have been studied extensively. However, it remains relatively understudied for Bluetooth low energy (BLE), especially in densely-deployed scenarios. In this work, we develop a testbed to conduct our experimental studies, focusing on BLE and its coexistence capabilities when being deployed in a dense environment, under possible interference from WiFi and ZigBee/IEEE 802.15.4. One scenario of interest is a network of several co-located BLE-based BANs, each of which is designed in a star topology with one gateway and multiple BLE sensor nodes. The second scenario represents a highly heterogeneous network where each BAN now carries both BLE and ZigBee sensors, while being exposed to interference from external WiFi transmission. Experiments are carried out on our testbed, which are built based on low-cost, light-weight off-the-shelf components and state-of-the-art BAN protocols. Our results show that the performance of BLE is relatively robust to interference from other BLE transmissions as well as those from nearby ZigBee and WiFi devices. In addition, the deployment of the testbed on human bodies results in no performance degradation for the network.**

***Index Terms*—Body area network, Bluetooth low energy, wearable sensors, coexistence, energy efficiency.**

I. Introduction

WIDESPREAD development of the Internet-of-Things (IoT) and low-power wireless technologies has led to profound interest in wireless body area networks (BANs). A BAN is composed of multiple nodes of very low-power, short-range sensors, sending and receiving data through wireless technologies. The sensors collect data from their surrounding environments, such as temperature, pressure, humidity or physiological conditions from a human body. Primary applications of BANs are in the healthcare domain, e.g., for ubiquitous monitoring of patients with chronic disease such as heart attack; but other use cases include military and sports.

Existing wireless technologies suitable for BANs are Bluetooth, Bluetooth low energy (BLE), or ZigBee/IEEE 802.15.4, which all operate in short range with low power. BLE is a widely adopted technology for short-range communication, developed in the distinctive feature of Bluetooth 4.x specification [2], [3]. Compared to the classic Bluetooth, it offers considerably reduced power consumption. It can also support higher data rate and lower latency than ZigBee [4], [5]. Thus, BLE is a very promising technology for BANs.

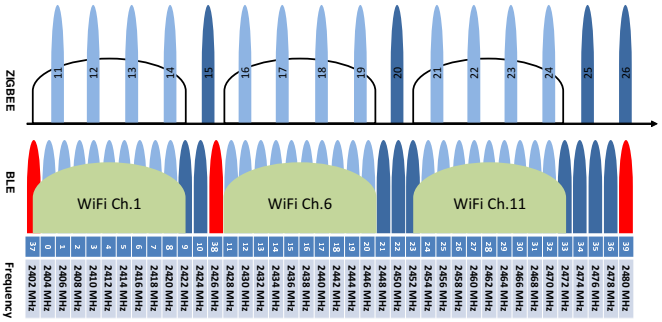

Fig. 1. Spectra of multiple wireless technologies in the 2.4-GHz ISM band. The colors indicate: *light-green* for WiFi bands; *red* for BLE's advertising channels 37, 38 and 39; *light-blue* for BLE's and ZigBee's data channels which overlap with WiFi channels; and *dark-blue* for BLE's and ZigBee's data channels which do not overlap with WiFi channels (and are free from WiFi interferences).

The heterogeneous candidate technologies for BANs above share the same frequency bands, i.e., the 2.4-GHz industrial, scientific, and medical (ISM) radio bands, which is notably home to not only Bluetooth, BLE, ZigBee but also WiFi. As shown in Fig. 1, BLE uses 40 narrow-band channels of 2 MHz bandwidth (advertising channels 37–39 and data channels 0–36). ZigBee also has narrow-band channels of 2 MHz bandwidth with a total of 16 channels spaced by 5 MHz. The WiFi channels are 20 MHz wide; and channels 1, 6 and 11 are non-overlapping channels separated by 5 MHz and commonly used for transmission. Interferences among multiple wireless technologies will occur due to the crowded shared frequency band. The use of incompatible modulations and channel access schemes makes it difficult to guarantee performance of devices across different technologies. Therefore, it is of vital importance to study their coexistence by considering mutual and cross-technology interferences, especially for BLE.

A. Related Work

We look at previous literature on BLE-based BANs and their coexistence. First, there are some works on implementing BLE platforms to investigate their features and performances. A

The material in this paper was presented in part at IEEE GLOBECOM, Singapore, December 2017 [1]. The editor coordinating the review of this paper and approving it for publication was E. Ayanoglu. (*Corresponding author: Mao V. Ngo.*)

Q. D. La, D. Nguyen-Nam, and H. T. Hoang are with the Singapore University of Technology and Design, Singapore 487372 (e-mail: duy.q.la@ieee.org; {namduong_nguyen, trunghieu_hoang}@sutd.edu.sg).

M. V. Ngo is with the Singapore University of Technology and Design, Singapore 487372, and also with the Institute for Infocomm Research (I^2R), A*STAR, Singapore (e-mail: vanmao_ngo@mymail.sutd.edu.sg).

T. Q. S. Quek is with the Singapore University of Technology and Design, Singapore 487372, and also with the Department of Electronics Engineering, Kyung Hee University, Yongin-si, Gyeonggi-do, 17104, Korea (e-mail: tonyquek@sutd.edu.sg).

BLE platform for remote health monitoring and compatibility for electrocardiography (ECG) monitoring was implemented in [5]; and the results showed good potential for medical applications in terms of throughput, end-to-end delay, and packet error rate, while staying energy efficient. Paper [6] provided experimental data on power consumption of BLE compared to ZigBee and ANT protocols[1] in a cyclic sleep scenario. The authors reported the lowest power consumption for BLE compared to ZigBee and ANT. In [7], a BAN testbed with one BLE master and four BLE slaves was set up for experimental evaluation, but external interferences were not considered. Meanwhile, [8] conducted spectrum survey to investigate the characteristics of BLE system by measuring the BLE transmission failure probability at a sport facility, university food court, and hospital intensive care unit. The work, however, did not specify interference sources. Mikhaylov *et al.* developed a simulation tool for BLE performances, first in the literature, and proposed two mechanisms to decrease interferences between BLE nodes, capable of improving throughput and energy efficiency in multi-node environment [9], [10]. Results from aforementioned works confirmed that BLE is resilient to the presence of high interference.

Cross-technology interferences and coexistence issues for BLE are investigated with other wireless technologies in several previous works. Silva *et al.* [11] conducted interference tests in an anechoic chamber for a pair of BLE devices, with a single interferer (either WiFi, ZigBee or classic Bluetooth). The results suggested good coexistence between BLE and WiFi, but slightly more collisions were observed between BLE and ZigBee as well as BLE and classic Bluetooth. Siekkinen *et al.* [12] compared the performance of BLE to ZigBee, under WiFi interference. Their results suggested that a minimum distance of 1.5 meters from the WiFi source is sufficient to avoid almost completely the interferences. When the interferer is very close, packet success rate for BLE was about 60%, while this rate for a ZigBee sensor was about 35% under similar settings. Bronzi *et al.* [13] presented their testbed to study BLE in inter-vehicular communications, with one pair of BLE devices and three pairs of Raspberry Pis in WiFi mode occupying channels 1, 6 and 11. Also, Natarajan *et al.* [14] examined pairwise coexistence between IEEE 802.15.4, BLE and WiFi, in which each network consisted of a pair of transmitter and receiver. They used the TI SensorTag based on the CC2650 multi-standard 2.4-GHz wireless micro-controller that can operate in either IEEE 802.15.4 or BLE mode. For IEEE 802.11b interference, an ad hoc wireless network was set up between two laptops. In their results, BLE was affected more by IEEE 802.15.4 interference than vice versa, and was more resilient than IEEE 802.15.4 against IEEE 802.11 interference.

Most of the above works suggested that BLE coexists well with WiFi and ZigBee. However, it is noted that they only considered very few devices and did not qualify as a *dense deployment*. A dense scenario was considered in [15], where the BLE latency and energy consumption under mutual interference were studied for one BLE pair with up to 15 other BLE pairs as interferers. However, this work only focuses on the BLE device discovery phase and does not offer insights during the BLE data transfer phase. In addition, none of the aforementioned works takes into account the *effects of human bodies* on the performance of BLE in a BAN setting. In fact, although there has been a considerable amount of research done in measuring and quantifying the effect of human bodies for classic Bluetooth (e.g., [16]) and IEEE 802.15.4 (e.g., [17], [18]), to our knowledge there is no work addressing the on-body deployment of BLE from a coexistence point of view. Overall, it is these two following key issues, dense BLE deployment on on-body BLE deployment, that represent a gap in BLE coexistence research and through this work we wish to address.

[1] ANT protocol is a proprietary wireless networking protocol using the 2.4-GHz frequency band, primarily for sensor networks, developed by ANT Wireless (a division of the Dynastream Innovations company).

B. Contributions and Paper Organization

Coexistence in the 2.4-GHz ISM bands, especially for BLE, has attracted a lot of works recently, as highlighted in our literature survey. Nonetheless, the following two deployment scenarios are currently lacking in coexistence studies: 1) densely-deployed BANs based on mutual BLE piconets (independent Bluetooth-based star-topology networks); and 2) dense heterogeneous BANs where BLE devices and those of other technologies are deployed simultaneously in the same BAN. These particular use-cases have important implications in modern days, where people can carry multiple wearable computing devices, under similar or different wireless interfaces (e.g., patients carrying both ZigBee-based electrocardiogram (ECG) sensor and BLE-based smart-band tracker) and they can be co-located in a dense area (e.g., a hospital). Enabling coexistence and interoperability (via IPv6) of heterogeneous devices is thus necessary to realize the full potential of the IoT [19], making it imperative to investigate the above two scenarios. We note also that coexistence tests and analyses for multiple piconets have been widely reported for classic Bluetooth [20]–[23]; but parallel results are currently absent for BLE. Last but not least, existing work does not offer any insight regarding the effects of human bodies on the performance of BLE in BANs. Motivated by that, our work considers multiple BANs based on BLE and investigates both dense BLE-based and dense heterogeneous networks in terms of coexistence for BLE, where WiFi interference can also be present. We conduct extensive experiments via our developed testbed in order to gain insights into the above key challenges. The two dense BLE deployment scenarios above are tested altogether for a static setting and an on-body setting. Thus, the contributions of this paper are four-fold:

- Presenting our developed BAN testbed, including the custom BLE-mote as the IPv6 BLE-enabled sensor platform, which is compatible with IoT devices and applications.
- Evaluating the coexistence of densely-deployed BLE-based BANs, with and without WiFi interference.
- Evaluating the coexistence of BLE in dense heterogeneous BANs where BLE and ZigBee are present simultaneously, with and without WiFi interference.

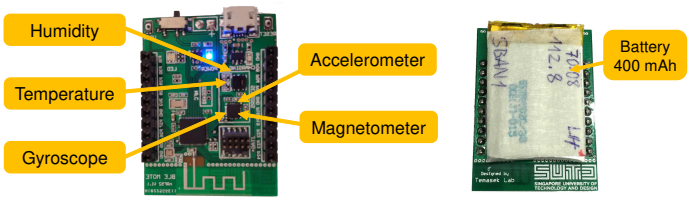

Fig. 2. The BLE-mote with (a) Front and (b) Back view.

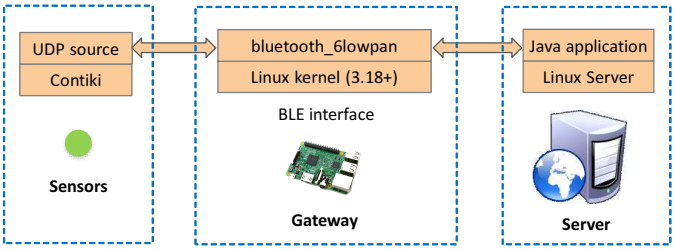

Fig. 3. The testbed's software specification.

- Evaluating the impacts of human body presence on the performance of BLE in densely deployed and dense heterogeneous BANs.

The rest of the paper is organized as follows. Section II introduces the testbed. Sections III and IV address experimental setup and results, respectively. Finally, Section V concludes this paper.

II. Description of the Testbed

A. General Topology

Our testbed consists of several independently deployed BANs in close proximity of each other. A single BAN consists of several wearable sensor nodes and a gateway connected in a star topology, which is intended for a single human body. A single BAN may contain sensors with highly heterogeneous functions for different monitoring purposes, including motion, ambiance, vital signs, and so on. The sensors, furthermore, mainly utilize BLE as their means of communications (and occasionally ZigBee). The sensor nodes will transmit their collected data to the gateway; and the gateway is responsible for either processing the data or forwarding the data to a common back-end server for further processing. We allow the gateways to use wired Internet connections to the server in order to focus only on the interferences across different BANs. As heterogeneous nodes content for the available wireless resources to transmit their data, we are interested in their performance under mutual and cross-technology interference.

B. Hardware Specification

1) BLE Platform – The BLE-mote: The BLE-mote is a BLE development kit which was developed at Temasek Labs, Singapore University of Technology and Design (SUTD) in 2016. Each BLE-mote is an integrated circuit (IC) board upon which several sensors can be mounted. Fig. 2 shows the layout of a BLE-mote. The BLE-mote board is itself based on Nordic Semiconductor's nRF52832 IC [24].

The BLE-mote has the following specifications:

a) Radio communications: 2.4-GHz transceiver; -96 dBm sensitivity in BLE mode; single-pin antenna interface; 1–2 Mbps supported data rates; Tx power: -20 to +4 dBm in 4 dB steps.
b) Microcontroller unit (MCU): ARM® Cortex®-M4 32-bit processor with floating-point unit (FPU), 64 MHz; 512 kB flash memory/64 kB RAM.
c) Built-in Sensors: MPU9250 [25] (triaxial accelerometer, magnetometer and gyroscope); and SHT21 [26] (humidity and temperature) sensor chips.
d) Battery: Polymer Lithium-Ion, 3.7 V at 400 mAh.

The novel features of the BLE-mote are that it is one of the first development kits that supports BLE ver. 4.2 and IPv6 capability as an enabler of the IoT (as of Mar. 2017); and it allows for easy integration of a wide range of off-the-shelf sensor and healthcare devices.

2) Gateway: We choose Raspberry Pi 3 (RPi3) [27] as the hardware device for the gateway. RPi3 is an off-the-shelf single-board computer which supports BLE 4.1. One of the main advantages of using RPi3 in our testbed is that it allows for adoption of the latest open-source BLE stack in Linux kernel, i.e., the BLE IPv6 over low power wireless personal area networks (6LoWPAN) gateway module. Moreover, this solution follows the IoT standards (e.g., IPv6 over BLE [28]), which will enable interoperability of the platform with other off-the-shelf IoT products.

C. Software Specification

The diagram in Fig. 3 shows the software interfaces between the different devices within the testbed. For the sensor platform (BLE-mote), we use the Contiki operating system (OS) [29], an open source OS for the IoT. Contiki is implemented using the software development kit (SDK) provided directly by Nordic. On the application layer, the BLE-motes run a user datagram protocol (UDP)-sender application to send the packets to the server via the gateway. We use the IPv6 over BLE protocol stack [28], with IPv6 layer on top of BLE 6LoWPAN and BLE logical link control and adaptation protocol (BLE L2CAP).

On the other hand, the BLE interface of the RPi3 gateway communicates with the BLE-motes directly via an open-source module called `bluetooth_6LoWPAN` on top of Linux Kernel (version 3.18+). RPi3 then also forwards packets to the remote server, via an Ethernet interface.

We also utilize *Collector-View*, a two-part application modified from an open-source framework provided along with Contiki-OS for our purposes. The first part of *Collector-View* runs on sensor platforms, which is implemented in C and serves to forward measurement data to gateway and server. The second part, implemented in Java, runs on the server. It collects sensor data via UDP connection through RPi3 gateway, then visualizes data graphically for analytical purposes.

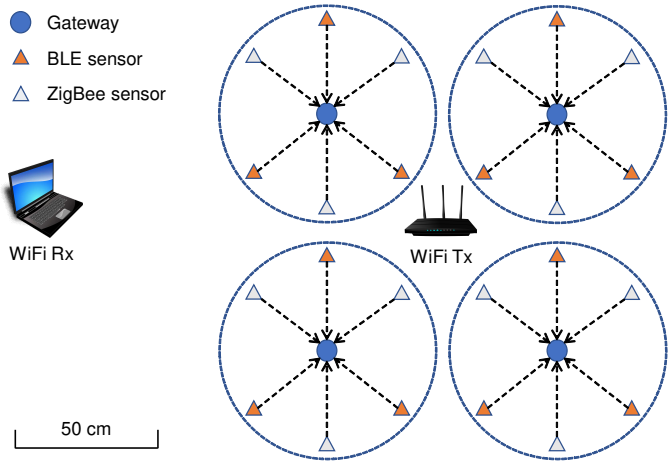

Fig. 4. Schematic of the testbed placement.

III. Experimental Setup

A. Scenarios

Our experiments address two main BLE coexistence issues for 1) dense BLE-based BANs; and 2) dense heterogeneous BANs. As such, the following two scenarios are considered:

1) Dense BLE-based BANs: In this scenario, each BAN is an independent BLE piconet, i.e., a star network with a gateway at the center and 3 surrounding BLE-motes at an approximately 0.3 m distance from the gateway. The number of BANs is varied from 1 to 4. Neighboring BANs are placed within a 0.5 to 2 m distance from each other (measured from the gateway). The whole testbed is placed inside an indoor laboratory and on an even surface (see Fig. 4).

2) Dense heterogeneous BANs: In this scenario, the testbed placement is similar to the previous one and the number of BANs is also varied from 1 to 4. However, each BAN now not only contains 3 BLE nodes but also 3 other ZigBee nodes simultaneously. The 6 nodes also are placed around a common gateway in a star topology at around 0.3 m distance. Therefore, at maximum capacity, the system can contain up to 24 heterogeneous devices transmitting simultaneously all co-located within a small area. We use the OpenMote-CC2538 combined with the OpenBattery board [30] as the ZigBee transmitting device, on which various sensors have also been built in, e.g., SHT21 (temperature and humidity), MAX44009 (light), and ADXL346 (accelerometer). One OpenMote-CC2538 device is plugged into the RPi3 gateway via serial port acting as the sink node in order to collect data from other transmitting ZigBee nodes. Time-slotted channel hopping (TSCH) protocol [31] is used as the medium access control (MAC) protocol for ZigBee devices, which is configured to do frequency hopping over a 16-frequency sequence.

Fig. 4 shows a schematic of the testbed placement in the laboratory. For the first scenario, the ZigBee nodes are switched off. We first run the experiment with one BAN switched on, then gradually the other BANs can be added with up to 4 BANs simultaneously.

B. External Interference

Ambience: Our testbed is run under the normal everyday environment of the laboratory. No intentional interferer is present. Any ambience signals in the wireless medium (i.e., the 2.4-GHz band) come from neighboring wireless networks such as the office's wireless local area networks (WLANs). We refer to their presence as the ambience, which is present in all scenarios and considered a benchmark.

ZigBee/IEEE 802.15.4: In a heterogeneous BAN configuration such as the one used in our second scenario above, both BLE and ZigBee sensor nodes can be carried by the same BAN. However, in the scope of this work, we treat coexisting ZigBee nodes in the same BAN as interferers.

WiFi: A nearby WiFi interferer can be intentionally introduced in order to study the coexistence of our testbed with WiFi networks. We set up a private WLAN through which a laptop downloads a large file from a wireless router. The interferer uses WiFi channel 1 throughout the experiments and the speed limit for downloading is set at 32 Mbps. The locations of the WiFi transmitter and receiver relative to the BANs are also depicted in the diagram in Fig. 4. This setup is applied to both scenarios in Section III-A.

C. Measurement Approach

1) Power Consumption: This directly measures power efficiency, defined as the averaged power consumption in mW of a BLE-mote across the duration of one BLE connection event. Fig. 5 shows our current measurement using a 10 Ω shunt resistor during a BLE connection event, whose profile perfectly matches with the one by Nordic [24]. The major activities take place within a very short period of roughly 2.4 ms during a connection interval (67.5 ms). The main states of a BLE-mote can be roughly divided into: CPU-active state (A, B, D, E, F, G), low-power state (C, H), Radio-Rx state (E) and Radio-Tx state (F, first half of G). The states can be overlapping but CPU-active and low-power states are mutually exclusive and altogether make up an entire interval. The average power can be computed from the currents, voltage and the measured time ticks in software, as follows:

$$P = \frac{V(I_{CPU}T_{CPU} + I_{LP}T_{LP} + I_{Rx}T_{Rx} + I_{Tx}T_{Tx})}{T_{CPU} + T_{LP}}. \quad (1)$$

where T_{CPU}, T_{LP}, T_{Rx}, T_{Tx}, I_{CPU}, I_{LP}, I_{Rx} and I_{Tx} denote the total times spent as well as the average currents in the four above states, respectively; and V denotes the voltage. Specifically, from Nordic's profiler tool [32], $I_{CPU} = 3.712$ mA, $I_{LP} = 2$ μA, $I_{Rx} = 6.0$ mA, and $I_{Tx} = 6.2$ mA (these are average values). Moreover, $V = 3$ V.

2) Radio Duty Cycle (RDC): RDC measures the percentage of time the radio module of the device is on, in percentage (%). RDC enables fair comparison of protocols across hardware platforms [19], which may have different clocking profiles but similar timing of radio transmission. It is given by

$$\text{RDC} = \frac{T_{Rx} + T_{Tx}}{T_{CPU} + T_{LP}}. \quad (2)$$

3) UDP Packet Delivery Ratio (PDR): We use the UDP packet delivery ratio as a metric for reliability. PDR, measured

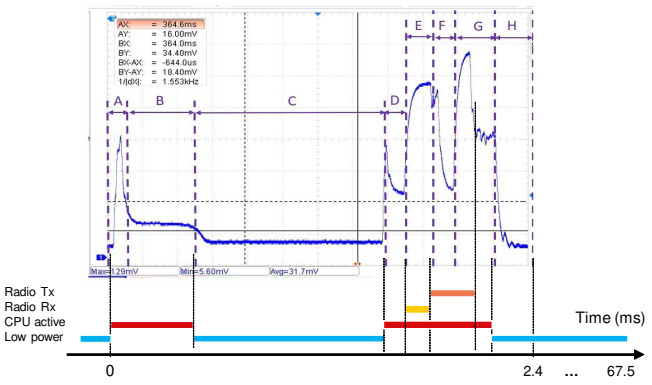

Fig. 5. The power consumption profile of BLE.

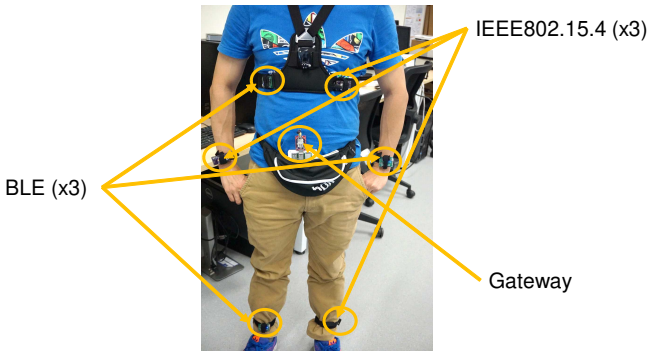

Fig. 6. On-body BAN deployment.

in percentage (%), is defined as the percentage of UDP packets successfully delivered from the source (BLE-mote) to the destination (server). Specifically, we count number of UDP packets received at the server and track the consecutive sequence number of these packets to capture missing packets over time. Let $\mathsf{N}_{\mathrm{UDP}}$ and $\mathsf{N}^{*}_{\mathrm{UDP}}$ be the total number of transmitted UDP packets and successful UDP packets, respectively. Then, the PDR is given by

$$\mathrm{PDR} = \frac{\mathsf{N}^{*}_{\mathrm{UDP}}}{\mathsf{N}_{\mathrm{UDP}}}. \tag{3}$$

4) BLE Packet Reception Ratio (PRR): PRR is a measure of one-hop communication between the BLE-mote and the gateway from a link-layer perspective. PRR is defined as the percentage of BLE packets successfully received from a BLE-mote at the RPi3. The number of BLE packets is not the same as the number of UDP packets; thus PRR and PDR are two different measures. The reasons for having these two separate ratios are not only layer-dependent but also to account for one-hop vs. end-to-end communications (e.g., two separate ratios are considered in [33]). Also, BLE retransmissions are not reflected in PDR at higher layer but can be captured by PRR; hence PRR usually has smaller values. Similarly, let $\mathsf{N}_{\mathrm{BLE}}$ and $\mathsf{N}^{*}_{\mathrm{BLE}}$ be the total number of transmitted BLE packets by BLE-motes and the number of successful BLE packets at the respective gateways, respectively. Then, the PRR is given by

$$\mathrm{PRR} = \frac{\mathsf{N}^{*}_{\mathrm{BLE}}}{\mathsf{N}_{\mathrm{BLE}}}. \tag{4}$$

D. Parameters

For each combination of network topologies and interferers, an experiment lasts for *two hours*. The various measurement metrics of interest are recorded and logged at the server. At the end of an experiment, statistics such as power consumption, RDC, PDR for UDP packets and PRR for BLE packets are averaged over all individual nodes. For all the nodes, transmission power is set at *0 dBm*.

The amount of data to be collected by nodes and transferred back to the server is usually application-dependent. In this work, we consider data from accelerometer sensors, which can have applications such as in healthcare and patient activity monitoring. Thus, the data payload consists of samples from the triaxial accelerometer of each sensor node. On each accelerometer, 6 bytes of data are collected per sample (3 axes $\times$ 2 bytes/axis). With the default sensor's sampling rate of 32 Hz, this amounts to *192 bytes of payload* generated per second. The *sending rate* (i.e., rate at which new UDP packets are sent to the gateway by sensors) is set at *6 Hz*. That is, each UDP packet contains $\frac{192}{6} = 32$ bytes of application data.

Per the IPv6 over BLE protocol stack [28] and the 6LoWPAN specification [34], apart from application data, a UDP packet also contains LoWPAN header (1 byte), IPv6 header (40 bytes), and UDP header (9 bytes). This adds up to a total of 49 bytes. With the Nordic BLE stacks, each BLE packet can carry only 27 bytes of data. Therefore, it takes $\lceil \frac{49+32}{27} \rceil = 3$ BLE packets to transmit one UDP packet; or *18 BLE packets per second*. For ZigBee, the 32-byte data, plus headers, can fit into a single IEEE 802.15.4 packet (at PHY/MAC layer), which can contain up to 127 bytes of data. Therefore, at 6 Hz sending rate, 6 IEEE 802.15.4 packets per second can be sent to the gateway.

E. On-body Deployment

The setup described above is a static placement where the BANs are immobile and are deployed off the human bodies, for ease of conducting the experiments. Practically, however, IoT devices are carried around by human users, which might be subjected to more randomness due to the presence of the human bodies and their movements. Thus, we also repeat the whole experiments (including two original scenarios with different combinations of external interferers) for the case of on-body BAN deployment.

Participated in the experiments are four male adults of medium builds, who wear a BAN set (3 BLE nodes, 3 IEEE 802.15.4 nodes, and a gateway) on their bodies. Fig. 6 shows the deployment of the BAN on a human body. The sensor nodes' positions are: left and right wrists, left and right ankles, and left and right sides of chest, as shown in the figure. The gateway is worn around the user's waist in this study.

The four users wearing BANs are seated in the same positions as highlighted in the schematic shown in Fig. 4. During all the experiments, the users carry out normal office activities, i.e., sitting and working in front of laptops and workstations,

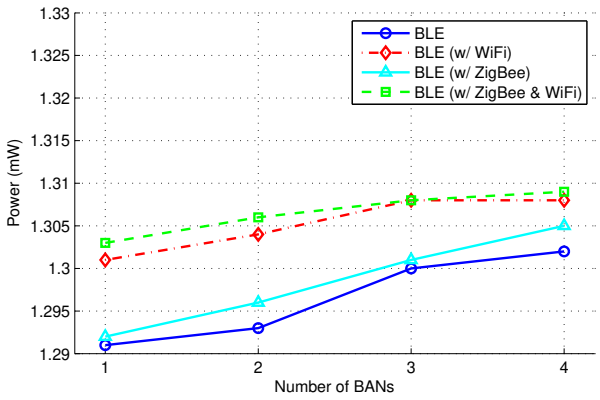

Fig. 7. Average power consumption of BLE nodes vs. the number of BANs.

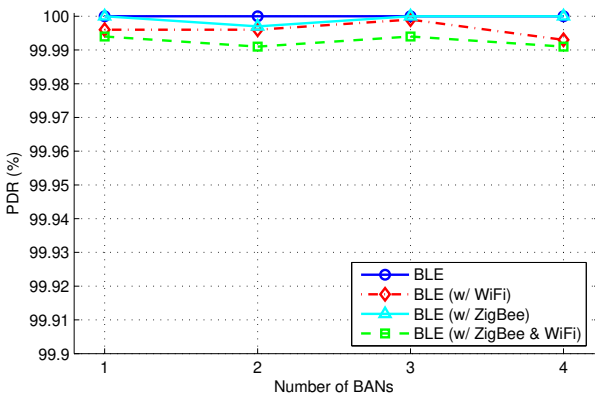

Fig. 9. Average UDP packet delivery ratio vs. the number of BANs.

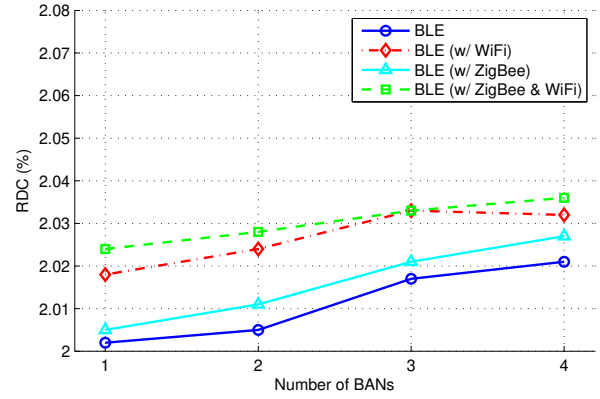

Fig. 8. Average radio duty cycle of BLE nodes vs. the number of BANs.

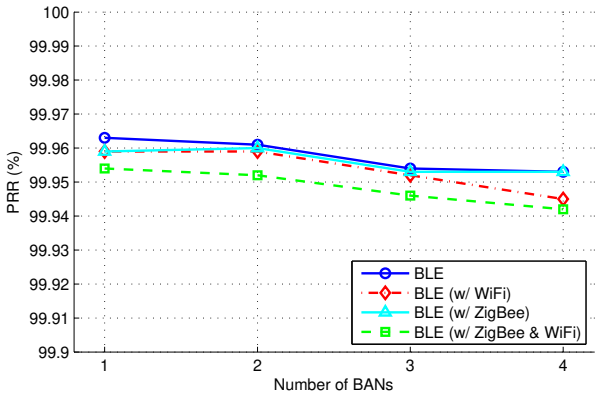

Fig. 10. Average BLE packet reception ratio vs. the number of BANs.

with natural hand, foot and body gestures. Sudden and high-mobility movements such as standing up and walking are allowed but limited to a minimum throughout the experiments; and users are to stay within a 3×3 m^2 area.

The parameter settings for on-body BAN deployment follow the previous setup for static experiments. However, experiments are only conducted in *fifteen-minute intervals* to minimize discomfort on the human users.

IV. Experimental Results

Using the measurement approach and the testbed setup described above, we perform our coexistence experiments for each scenario. We allow the number of BANs to vary from 1 to 4; and each configuration is subject to both with and without WiFi interference. Subsequently, we report our measurement results in Section IV-A, which applies to the static placement. In addition, Section IV-B provides additional insights when the BANs are deployed on the human bodies.

A. Measurement Results

The results of the coexistence experiments are shown in Figs. 7 to 10. We look at each individual performance metric as follows.

1) Power Consumption: The measurement results for average power consumption of BLE nodes are collected and displayed in Fig. 7 against the number of BANs. The results are further divided into 4 groups: BLE-based BANs, with and without WiFi interference; and BLE in heterogeneous BANs (with ZigBee), with and without WiFi interference. They are abbreviated by *BLE*, *BLE (w/ WiFi)*, *BLE (w/ ZigBee)* and *BLE (w/ ZigBee & WiFi)*, respectively.

Our first observation is that as the number of BANs increases, power consumption also grows in all groups, although the margin is relatively small. On average, each BLE node in the 1-BAN ambient case consumes about 1.291 mW. When the number of BANs goes up to 4, this figure only increases to 1.302 mW, i.e., by 0.85%. Similar figures can be seen with the other cases. It is thus implied that when the network is denser, nodes should engage more in radio transmission, possibly as more collisions and retransmissions occur, causing power consumption to increase (but not much). We should therefore see the same trend in RDC in Fig. 8, which will be discussed later.

Regarding the impacts of interferers on BLE power consumption, the results in Fig. 7 seem to indicate that both ZigBee and WiFi presences cause BLE to spend more power; and the effect of WiFi interference is more apparent. In fact, for the dense BLE-piconet BANs (marked by circles and dia-

monds), introducing WiFi causes a 0.009 mW power jump on average; and for the heterogeneous BANs (marked by triangles and squares), 0.008 mW. Meanwhile, by introducing ZigBee simultaneously with BLE, the above figures are reduced to 0.002 mW and 0.001 mW, respectively. Nevertheless, the amount of average power increment is relatively tiny. Thus, our data, to some extent, have verified that BLE is relatively resilient to dense deployment as well as cross-technology interference, in terms of power consumption.

To further put things into perspective, we also compare the average BLE power with that of ZigBee nodes, which is around 6.093 mW (not shown in the figures), i.e., 4.7 times higher than BLE's. This suggests that BLE is highly more power efficient compared to ZigBee. The result is in line with previous findings which showed a superior energy efficiency of BLE over ZigBee [6], [12].

2) Radio Duty Cycle: The measurement results for RDC are displayed in Fig. 8. The grouping for RDC is similar to those of power consumption in the previous section. We can see from Fig. 8 that RDC shows a remarkably strong correlation to power consumption, as is expected. RDC also grows slightly with the number of BANs; and introducing interference, especially WiFi, also accounts for a higher percentage in the duty cycle. This is also due to the growth in collisions and retransmissions, leading to more time spent in radio modes, similar to the one observed in power consumption. We should note that the average percentage of RDC in all cases stays roughly above 2% which takes a very small proportion of one BLE connection event. Our RDC results could be compared to those of [19] where RDCs of 0.5-1.3% were recorded for request-response mode and 28.8% for bulk transmission mode. In contrast, our scenarios could be categorized as *periodic traffic* mode due to the constant stream of sensor data, which should lie in between the two cases above.

3) UDP Packet Delivery Ratio: Fig. 9 displays the PDR versus the number of BANs. Overall, PDR measurements are consistently high in all cases, at above 99.99%, and even at 100% for BLE without interferers. This shows that for BLE, delivery of the application payloads can be done reliably even in the presence of dense neighboring devices as well as various cross-technology interference. This might be attributed to the BLE retransmission mechanism, so that even if a collision occurs at the link layer and a BLE packet is lost, it is still possible for the entire UDP packet to be successfully received at the server. The presence of WiFi causes some slight drop in PDR ($< 0.01\%$) over the 2-hour experiments, which is relatively insignificant and comparable to results in [13] (in [13] only 1-minute experiments were reported).

4) BLE Packet Reception Ratio: Fig. 10 displays the PRR results. It first suggests that there are more collisions and packet losses in the link layer than for UDP packets (although the success ratios are still above 99.9%). Collisions seem to occur more often as the spectrum gets congested due to increased number of interfering devices but the PRR degradation is not overall severe. We note that at our sending rate of 18 BLE packets/s per each BLE-mote, the results show marked PRR improvement over equivalent results [5]. Also, the separation among multiple test scenarios is not clear-cut, confirming the BLE performance robustness in crowded scenarios, with and without cross-technology interference.

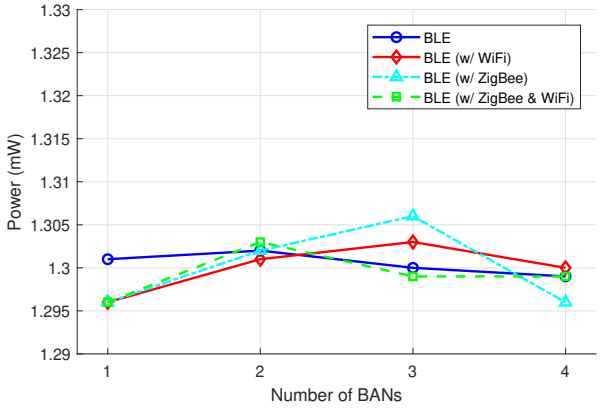

Fig. 11. Average power consumption vs. the number of BANs for the case of on-body deployment.

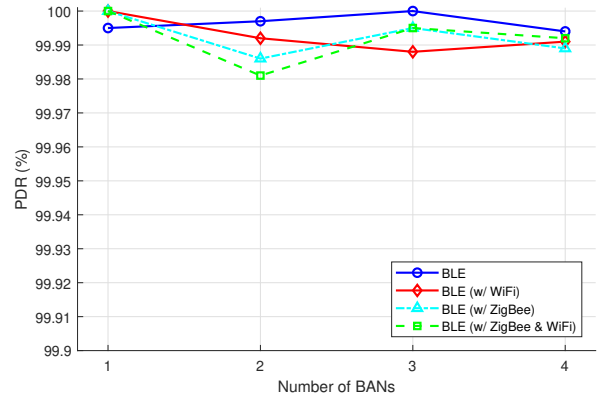

Fig. 12. Average UDP packet delivery ratio vs. the number of BANs for the case of on-body deployment.

B. Effects of the Human Bodies

We next address the impacts of on-body BLE deployment in dense scenarios on the performance of the network. Without loss of generality, we look at two representative metrics which are power consumption and PDR. Figs. 11 and 12 display the measured power consumption and PDR, respectively. We notice two key observations: 1) there is *no significance difference* in terms of the range of power and PDR obtained, as compared to Figs. 7 and 9; and 2) there are *more fluctuations* in the trends of power and PDR as the number of BANs increases from 1 to 4 for both metrics.

Evidently from Fig. 11, introducing the human users does not cause any observable deviation in the power consumption of the BLE nodes from previously recorded results. In fact, the measurements stay consistently within the 1.290-1.310 mW range, which is similar to those in Fig. 7. The overall mean power consumption of a BLE node, averaged over all 24 separate experiments, is about 1.302 mW for the static case, while it is 1.300 mW for the on-body case. This is inconclusive of any significant impacts from the human users on the energy

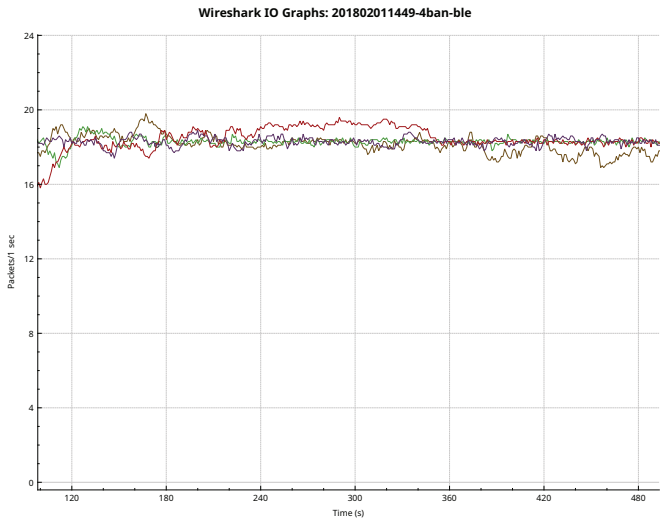

Fig. 13. Wireshark capture of the number of packets received over time for 4 on-body BANs with ambience interference.

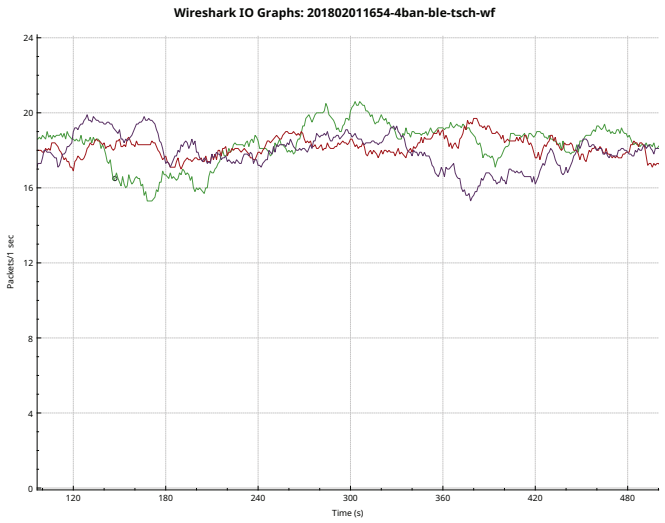

Fig. 14. Wireshark capture of the number of packets received over time for 4 on-body BANs with ZigBee and WiFi interferences.

usage of nodes. On the other hand, while Fig. 11 shows a slight hint of increasing power as more BANs are simultaneously present, this tendency is masked by fluctuations and anomalies (e.g., the 4-BAN BLE (w/ ZigBee) measurement). If anything, this suggests to some extent that the presence of human bodies and their movements did cause fluctuation in the network performance compared to the off-body results, but far from being detrimental.

The PDR results in Fig. 12 supports this claim further. We see that the PDR results for the on-body experiments also lie in the high end of the 99.90-100% range, similar to those in Fig. 9. At the same time, the fluctuations seem to be more visible. Nevertheless, such a high PDR result indicates that even with the human bodies, almost all the UDP packets are successfully received at the destinations.

We are yet to see the effects of different types of interferers in the on-body deployment scenarios, possibly because the fluctuation observed above plays a dominant role. To gain further insights into this, we run the packet analyzer Wireshark at the server to capture real-time UDP packet reception at the end destination. Figs. 13 and 14 represent the data for two illustrative cases, i.e., the 4 on-body BLE-based BANs without interference, and 4 on-body heterogeneous BANs with WiFi interference. In each figure, four lines can be seen, corresponding to UDP packet streams from four different gateways. Ideally, due to our settings, each BAN wearer should generate 18 UDP packets per second (6 Hz UDP sending rate/node $\times$ 3 nodes/BAN). It can indeed be seen from the two figures that all streams are positioned around the 18 packets/s mark, subject to some degree of fluctuations. Under interferences from ZigBee and WiFi (Fig. 14), the gaps between different lines become wider and more uneven, suggesting that interferences can still cause packet loss to occur more frequently at a given point in time. However, BLE retransmission is able to maintain the overall success rate for UDP packets, thus keeping the PDR at a high level regardless of external interferences and human presence.

In general, it is fair to say that BLE remains power efficient and resilient to interferences even when deployed on human bodies. It should also be pointed out that compared to the 2-hour static experiments, the on-body ones are conducted for 15-minute periods so a single packet loss might cause a much greater drop on the PDR than earlier, which could add to the fluctuation. Unlike the off-body measurement results in Section IV-A where we could make comparison with equivalent results in the literature, it is difficult to do so here due to the lack of work on coexistence of on-body BLE. Nevertheless, it is interesting to note that two of our observations, i.e., that BLE is reliable and consistent despite on-body effects, and that BLE retransmission helps to deliver very high PDR, echoing those made in Shah *et al.* [16], albeit for classic Bluetooth.

V. Conclusions

In this work, we developed our own testbed to study the coexistence of BLE-based BANs in several densely-deployed scenarios, subject to possible cross-technology interference. We also considered a heterogeneous BAN setup with ZigBee nodes and BLE nodes in the same BAN. Our testbed experimental results suggested good coexistence capability of BLE, not only in a dense network of mutual BLE-based BANs, but also in a highly heterogeneous setting, under WiFi interference. When deploying BANs on human bodies, visible fluctuations in power consumption and PDR can be observed; however, it does not cause any noticeable change in the overall energy efficiency and reliability of BLE.

Acknowledgement

We would like to thank Dr. Thang V. Nguyen and Tuan Anh Bui for their help with the experiments; and Dr. Na-Rae Kim for her assistance with the preparation of the manuscript. Also, we thank the anonymous reviewers for their valuable comments and suggestions.

REFERENCES

[1] Q. D. La, D. Nguyen-Nam, M. V. Ngo, and T. Q. S. Quek, "Coexistence evaluation of densely deployed BLE-based body area networks," in *Proc. IEEE Globecom*, Singapore, Dec. 2017.

[2] Bluetooth Specification Version 4.2. [Online]. Available: www.bluetooth.org/DocMan/handlers/DownloadDoc.ashx?doc_id=286439

[3] R. Tabish, A. B. Mnaouer, F. Touati, and A. M. Ghaleb, "A comparative analysis of BLE and 6LoWPAN for U-HealthCare applications," in *Proc. IEEE-GCC Conference and Exhibition (GCCCE)*, 2013, pp. 286–291.

[4] E. Georgakakis, S. A. Nikolidakis, D. D. Vergados, and C. Douligeris, "An analysis of Bluetooth, Zigbee and Bluetooth Low Energy and their use in WBANs," in *Wireless Mobile Communication and Healthcare*, J. C. Lin and K. S. Nikita, Eds. Springer Verlag, 2011, pp. 168–175.

[5] F. Touati, O. Erdene-Ochir, W. Mehmood, A. Hassan, A. B. Mnaouer, B. Gaabab, M. F. A. Rasid, and L. Khriji, "An Experimental Performance Evaluation and Compatibility Study of the Bluetooth Low Energy Based Platform for ECG Monitoring in WBANs," *Int. J. Distrib. Sen. Netw.*, pp. 1–12, Jan. 2015.

[6] A. Dementyev, S. Hodges, S. Taylor, and J. Smith, "Power consumption analysis of Bluetooth Low Energy, ZigBee and ANT sensor nodes in a cyclic sleep scenario," in *Proc. IEEE IWS*, 2013, pp. 1–4.

[7] J. A. Afonso, A. J. Maio, and R. Simoes, "Performance Evaluation of Bluetooth Low Energy for High Data Rate Body Area Networks," *Wirel. Pers. Commun.*, vol. 90, no. 1, pp. 121–141, Sep. 2016.

[8] M. O. A. Kalaa, W. Balid, N. Bitar, and H. H. Refai, "Evaluating Bluetooth Low Energy in realistic wireless environments," in *Proc. IEEE WCNC*, 2016, pp. 1–6.

[9] K. Mikhaylov, "Simulation of network-level performance for Bluetooth Low Energy," in *Proc. IEEE PIMRC*, 2014, pp. 1259–1263.

[10] K. Mikhaylov and T. Hanninen, "Mechanisms for Improving Throughput and Energy Efficiency of Bluetooth Low Energy for Multi Node Environment," *Journal of High Speed Networks*, vol. 21, no. 3, pp. 165–180, Aug. 2015.

[11] S. Silva, S. Soares, T. Fernandes, A. Valente, and A. Moreira, "Coexistence and interference tests on a Bluetooth Low Energy front-end," in *Proc. IEEE Sci. Info. Conf. (SAI)*, 2014, pp. 1014–1018.

[12] M. Siekkinen, M. Hiienkari, J. K. Nurminen, and J. Nieminen, "How low energy is Bluetooth Low Energy? comparative measurements with Zigbee/802.15.4," in *Proc. IEEE WCNC Workshops*, 2012, pp. 232–237.

[13] W. Bronzi, R. Frank, G. Castignani, and T. Engel, "Bluetooth Low Energy performance and robustness analysis for Inter-Vehicular Communications," *Ad Hoc Netw.*, vol. 37, no. P1, pp. 76–86, Feb. 2016.

[14] R. Natarajan, P. Zand, and M. Nabi, "Analysis of Coexistence between IEEE 802.15.4, BLE and IEEE 802.11 in the 2.4 GHz ISM Band," in *Proc. Annu. Conf. IEEE Ind. Electron. Soc.*, 2016, pp. 6025–6032.

[15] J. J. Treurniet, C. Sarkar, R. V. Prasad, and W. d. Boer, "Energy Consumption and Latency in BLE Devices under Mutual Interference: An Experimental Study," in *Proc. IEEE Future Internet of Things and Cloud (FiCloud)*, 2015, pp. 333–340.

[16] R. C. Shah, L. Nachman, and C.-y. Wan, "On the Performance of Bluetooth and IEEE 802.15.4 Radios in a Body Area Network," in *Proc. ICST 3rd Intl. Conf. Body Area Networks (BodyNets '08)*, Mar. 2008, pp. 25:1–25:9.

[17] B. de Silva, A. Natarajan, and M. Motani, "Inter-user interference in body sensor networks: Preliminary investigation and an infrastructure-based solution," in *Proc. 6th Intl. Workshop on Wearable and Implantable Body Sensor Networks*, Berkeley, CA, USA, Jun. 2009, pp. 35–40.

[18] N. Amini, W. Xu, Z. Li, M.-C. Huang, and M. Sarrafzadeh, "Experimental analysis of IEEE 802.15.4 for on/off body communications," in *Proc. IEEE PIMRC*, Toronto, ON, Canada, Sep. 2011, pp. 2138–2142.

[19] P. Narendra, S. Duquennoy, and T. Voigt, "BLE and IEEE 802.15.4 in the IoT: Evaluation and Interoperability Considerations," in *Internet of Things. IoT Infrastructures: Second International Summit, IoT 360° 2015, Rome, Italy, October 27-29, 2015, Revised Selected Papers, Part II*, B. Mandler, *et al.*, Ed. Springer, 2016, pp. 427–438.

[20] A. Arumugam, A. Nix, P. Fletcher, S. Armour, and B. Lee, "Scenario driven evaluation and interference mitigation proposals for Bluetooth and high data rate Bluetooth enabled consumer electronic devices," *IEEE Trans. Consum. Electron.*, vol. 48, no. 3, pp. 754–764, Aug. 2002.

[21] I. Howitt, "Mutual interference between independent Bluetooth piconets," *IEEE Trans. Veh. Technol.*, vol. 52, no. 3, pp. 708–718, May 2003.

[22] T.-Y. Lin and Y.-C. Tseng, "Collision analysis for a multi-Bluetooth picocells environment," *IEEE Commun. Lett.*, vol. 7, no. 10, pp. 475–477, 2003.

[23] F. Mazzenga, D. Cassioli, P. Loreti, and F. Vatalaro, "Evaluation of packet loss probability in Bluetooth networks," in *Proc. IEEE ICC*, vol. 1, 2002, pp. 313–317.

[24] Nordic Semiconductor. nRF52832 product specification v1.1. [Online]. Available: http://infocenter.nordicsemi.com/pdf/nRF52832_PS_v1.1.pdf

[25] InvenSense Inc. MPU-9250 Product Specification Revision 1.0. [Online]. Available: http://store.invensense.com/datasheets/invensense/MPU9250REV1.0.pdf

[26] Sensirion. Datasheet SHT21. [Online]. Available: www.sensirion.com/fileadmin/user_upload/customers/sensirion/Dokumente/2_Humidity_Sensors/Sensirion_Humidity_Sensors_SHT21_Datasheet_V4.pdf

[27] Raspberry Pi (Trading) Ltd. Raspberry pi hardware. [Online]. Available: http://www.raspberrypi.org/documentation/hardware/raspberrypi/

[28] Internet Engineering Task Force (IETF). (2015) IPv6 over BLUETOOTH(R) Low Energy. [Online]. Available: http://tools.ietf.org/pdf/rfc7668

[29] Contiki: The Open Source OS for the Internet of Things. [Online]. Available: http://contiki-os.org/

[30] OpenMote-CC2538 platform. [Online]. Available: www.openmote.com/

[31] IEEE Computer Society, "IEEE Standard for Low-Rate Wireless Networks," *IEEE Std. 802.15.4-2015 (Revision of IEEE Std 802.15.4-2011)*, pp. 1–709, Apr. 2016.

[32] Nordic developer zone. nRF52 online power profiler. Accessed: 25-Mar-2017. [Online]. Available: http://devzone.nordicsemi.com/power/

[33] J. Ko, J. Eriksson, N. Tsiftes, S. Dawson-Haggerty, A. Terzis, A. Dunkels, and D. Culle, "ContikiRPL and TinyRPL: Happy Together," in *Proc. Workshop on Extending the Internet to Low Power and Lossy Networks (IPSN 2011)*, 2011.

[34] Z. Shelby and C. Bormann, *6LoWPAN: The Wireless Embedded Internet*, ser. Wiley Series on Communications Networking & Distributed Systems. Chichester, UK: Wiley, 2011.

Quang Duy La (S'09-M'14) received the B.Eng. degree (with first-class honours) in 2008, and the Ph.D. degree under the Nanyang Presidents Research Scholarship in 2013, both in electrical and electronic engineering from Nanyang Technological University (NTU), Singapore. He was a Postdoctoral Research Fellow at Temasek Laboratories, Singapore University of Technology and Design (SUTD) from 2015 to 2018. He has participated in research projects on radio resource allocation, smart grid, body area networks and network security. His research interests lie in game-theoretic and distributed algorithms for next generation communications systems, network security, mobile edge computing, as well as machine learning and big data analytics for the Internet of Things.

Dr. La is a co-author of the monograph "Potential Game Theory: Applications in Radio Resource Allocation" published by Springer in 2016.

Duong Nguyen-Nam received the B.E. and M.E. degrees in Electrical and Electronics Engineering from Ho Chi Minh City University of Technology, Vietnam in 2010 and 2014, respectively. Currently, he is a research assistant in the Information Systems Technology and Design Pillar, Singapore University of Technology and Design. His current research interests include wireless communications and networks, network intelligence, internet-of-things, edge computing, and machine learning.

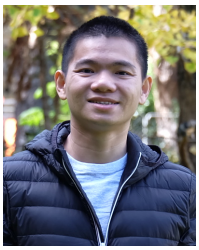

Mao V. Ngo (S'17) received the B.Eng. degrees (Hons.) in telecommunications engineering from the National Technical University of Ukraine, Kyiv Polytechnic Institute, in 2011, and the M.Eng. degree in electronics engineering from Ho Chi Minh City University of Technology, Vietnam, in 2015. He is currently pursuing the Ph.D. degree with the Information Systems Technology and Design Pillar, Singapore University of Technology and Design under the Singapore International Graduate Award (SINGA) A*STAR scholarship. His current research interests include edge computing, machine learning, wireless sensor networks, and scheduling for internet-of-things networks.

Hieu T. Hoang received the B.Eng. degree (Hons.) in electrical and electronics engineering (specializing in electronics), Ho Chi Minh City University of Technology, Vietnam, in 2016. He is currently working as a Research Assistant at Singapore University of Technology and Design (SUTD). His current research interests include edge computing, wireless sensor networks, machine learning, and Internet-of-things.

Tony Q.S. Quek (S'98-M'08-SM'12-F'18) received the B.E. and M.E. degrees in Electrical and Electronics Engineering from Tokyo Institute of Technology, respectively. At MIT, he earned the Ph.D. in Electrical Engineering and Computer Science. Currently, he is a tenured Associate Professor with the Singapore University of Technology and Design (SUTD). He also serves as the Acting Head of ISTD Pillar and the Deputy Director of the SUTD-ZJU IDEA. His current research topics include wireless communications and networking, internet-of-things, network intelligence, wireless security, and big data processing.

Dr. Quek has been actively involved in organizing and chairing sessions, and has served as a member of the Technical Program Committee as well as symposium chairs in a number of international conferences. He is currently an elected member of IEEE Signal Processing Society SPCOM Technical Committee. He was an Executive Editorial Committee Member for the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, an Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS, and an Editor for the IEEE WIRELESS COMMUNICATIONS LETTERS. He is a co-author of the book "Small Cell Networks: Deployment, PHY Techniques, and Resource Allocation" published by Cambridge University Press in 2013 and the book "Cloud Radio Access Networks: Principles, Technologies, and Applications" by Cambridge University Press in 2017.

Dr. Quek was honored with the 2008 Philip Yeo Prize for Outstanding Achievement in Research, the IEEE Globecom 2010 Best Paper Award, the 2012 IEEE William R. Bennett Prize, the 2015 SUTD Outstanding Education Awards – Excellence in Research, the 2016 IEEE Signal Processing Society Young Author Best Paper Award, the 2017 CTTC Early Achievement Award, the 2017 IEEE ComSoc AP Outstanding Paper Award, and the 2017 Clarivate Analytics Highly Cited Researcher. He is a Distinguished Lecturer of the IEEE Communications Society and a Fellow of IEEE.