Adaptive Transmission Power Control in BLE: Unveiling and Overcoming the Limits of Current Solutions

Elisabeth Salomon and Carlo Alberto Boano
Institute of Technical Informatics, Graz University of Technology, Austria
E-mail: {elisabeth.salomon, cboano}@tugraz.at

***Abstract*—Adapting the transmission (TX) power of Bluetooth Low Energy (BLE) devices at runtime is pivotal to maximize the reliability and efficiency of their communications. Unfortunately, research in this area is scant. Firstly, there is a notable gap in exploring the efficacy of utilizing received signal strength information to adjust the TX power, as mandated by BLE's newly-introduced LEPC feature. Secondly, the performance of existing approaches is constrained by their reliance on connection-wide information. In this work, we fill this gap and investigate in detail the performance of different schemes used to adapt the TX power of BLE devices at runtime. After shedding light on the suitability of different link quality indicators to inform the adaptation, we demonstrate experimentally that adjusting the TX power on a *per-channel* basis allows to drastically improve the performance of existing solutions. Finally, we propose and implement a novel TX power control mechanism for BLE devices embedding our empirical observations, showcasing its superior reliability and energy efficiency over traditional approaches.**

***Index Terms*—Adaptive protocols, Bluetooth Low Energy, LEPC feature, nRF52840, PRR, RSSI, Testbed experiments.**

I. Introduction

Due to its low power consumption and ubiquity in a vast array of consumer electronic devices, Bluetooth Low Energy (BLE) has become one of the most popular short-range wireless technologies to design Internet of Things (IoT) solutions. BLE radios are indeed not only pervasive in smartphones and wearables (*e.g.,* smart watches, fitness gadgets), but are often key enablers for location-based and navigation services (*e.g.,* keyless access [1] and indoor positioning systems [2]), as well as a number of use cases in safety-critical domains (*e.g.,* health monitoring [3], contact tracing [4], and asset tracking [5]).

Now at version 5.4, the Bluetooth Core specification has been regularly updated across the years to extend BLE's applicability to new use cases and improve its performance. For example, since the release of Bluetooth v5.0 in 2016, the core specification has been enriched with features such as direction finding, isochronous channels, as well as periodic advertisements with responses and encrypted advertising: these enable, respectively, better support for location-based services, low-power audio streaming, and secure bidirectional broadcast data exchange [6]. Extensions allowing to improve BLE's communication performance, include, among others, additional physical layers [7], channel map updates to deal with RF interference and changing environmental conditions, as well as dynamic transmission (TX) power control [8].

BLE's TX power control feature. To maximize the efficiency of BLE applications while coping with the dynamic nature of wireless links, the Bluetooth v5.2 standard introduced the so-called *Low Energy Power Control (LEPC)* feature. The latter can be used to adjust the TX power of a connected device based on received signal strength (RSSI) information. This feature is very important, as it allows a device to continuously find the right balance between communication reliability and energy efficiency. In fact, the use of a fixed TX power is often sub-optimal: conservatively choosing a high TX power upfront to account for potential link degradation would result in wasted energy. Conversely, aggressively opting for the lowest TX power to successfully communicate may result in poor reliability, especially in presence of device mobility, changes in environmental conditions, and fluctuations of the link quality. Moreover, the LEPC feature allows to avoid that a device uses an unnecessarily high TX power, which improves coexistence with co-located wireless devices operating in the 2.4 GHz band (*e.g.,* IEEE 802.15.4 and Wi-Fi devices).

The gap to fill. The Bluetooth standard only specifies that the LEPC feature should be used to keep the receiver's signal-to-noise ratio (SNR) in a *"golden range"*, and that the adaptation should be based on *reliable RSSI measurements*: the implementation of the actual algorithm and the definition of such golden range are left up to the developer. Unfortunately, research on TX power adaptation in BLE [9], [10] has mainly been carried out before the formalization of LEPC, and hence does not shed light on its effectiveness nor on how to maximize its performance. Moreover, related literature adapts the TX power based on the packet reception ratio (PRR), and does not consider RSSI information [9], [11]: it hence remains unclear how existing approaches compare to LEPC. Finally, existing solutions adjust the TX power equally for all channels and not *on a per-channel basis*. Since the quality of wireless links is strongly related to the employed frequency, adapting the TX power for each BLE channel separately may enhance performance [12]: this, however, has not been verified empirically.

Contributions. We fill this gap and study in detail the performance of different techniques used to adapt the TX power of BLE devices at runtime. We do so using experimental data collected on a real-world testbed in several environments (§ IV).

First, we study the performance of TX power adaptation based on RSSI, *i.e.,* the same information leveraged by Bluetooth v5.2's LEPC feature (§ V-A). We then show the superiority of *per-channel adaptation*, which shows a better trade-off between communication reliability and energy consumption (§ V-B). Second, we study the performance of TX power adaptation based on PRR, *i.e.,* the information used by state-of-the-art approaches (§ VI-A), showing that *per-channel adaptation* can reduce energy usage by up to 50% (§ VI-B).

We finally propose a novel TX power control mechanisms that combines both PRR and RSSI information together with per-channel adaptation and channel blacklisting, showing its superior reliability and efficiency over classical approaches (§ VII).

Before illustrating our results, we provide a brief overview of BLE technology (§ II) and review related literature (§ III).

II. BLE Fundamentals

BLE devices operate on 40 frequency channels within the license-free 2.4 GHz ISM band, and commonly transmit data at 1 Mbps. Since the release of Bluetooth v5.0 [13], devices can support additional PHY modes allowing to trade data throughput in favour of reliability and communication range [7]. Specifically, BLE devices can use either *coded* (125 and 500 kbps), which embed forward error correction techniques and incur a higher energy consumption, or *uncoded* PHYs (1 and 2 Mbps), which are faster albeit less reliable.

BLE supports *connection-oriented* and *connection-less* communication: the latter is commonly used for unidirectional data exchange[1], whereas the former allows for bidirectional flow of information between a central and a peripheral device. In connection-oriented communication, which is the focus of this work, link-layer acknowledgments are used to autonomously re-transmit packets and ensure a reliable delivery. Central and peripheral devices periodically exchange packets in so-called connection event (CEs), which are scheduled at fixed time intervals defined by the connection interval (CI). The CI hence directly affects the energy consumption of the devices[2], as it specifies how often they should be involved in transmission or reception activities. If there is no data to send in a CE, the device sends a link-layer keep-alive packet with empty payload to keep the connection active. BLE also allows a peripheral device to skip a specified number of CEs (dictated by the slave latency (SL)) in case there is no data to be sent: this allows a device to remain in a low-power state and save energy.

In connection-oriented BLE communication, a different channel is used in each CE[3]: this approach is called frequency hopping and allows to deal with unreliable channels. In fact, if a transmission fails on a given channel due to an insufficient link quality (*e.g.,* due to multi-path effects or cross-technology interference), future re-transmission attempts will be made in different channels, which typically helps solving frequency-dependent issues [14]. In case of persistently-unreliable channels, the BLE standard allows the exclusion of specific data channels (a.k.a. blacklisting) from the frequency hopping sequence using so-called *channel map updates*.

LE Power Control. Bluetooth v5.2 introduced the *Low Energy Power Control (LEPC)* feature for devices that have established a connection [8]. Devices supporting the LEPC feature coordinate to ensure that the current RSSI of incoming packets ($RSSI_{curr}$) falls into a *"golden range"* ($RSSI_{min} \leq RSSI_{curr} \leq RSSI_{max}$)[4]. Should the RSSI of received packets fall outside this range, a device can request the connected party to change its TX power by a specified delta in dB (at most ± 8 dB per request). To this end, a device can send specific link-layer control packets (`LL_POWER_CONTROL_REQ`[5]); both central and peripheral are allowed to initiate such requests. The request only affects a specific PHY mode and must be answered with a `LL_POWER_CONTROL_RSP`. A device should apply the proposed TX power adjustment before sending the response.

BLE devices are also allowed to autonomously change their TX power (at most once per second). In case they have previously received a `LL_POWER_CONTROL_REQ` packet, they should notify the corresponding sender about this change by transmitting a `LL_POWER_CHANGE_IND` packet.

III. Related Work

TX power adaptation is a topic that has been largely investigated in the context of wireless sensor networks and low-power IoT technologies over the last decade. The literature is dense of both theoretical and experimental studies on *short-range* wireless technologies such as IEEE 802.15.4 [15]–[19] and *long-range* wireless technologies such as LoRa [20], [21].

TX power control in BLE. When it comes to BLE systems, the body of research on TX power adaptation is relatively small. Most studies proposing dynamic adjustments of the TX power of BLE devices aim at increasing ranging performance [22]–[24]. In fact, one can leverage multiple TX power levels to better characterize the link between two devices through RSSI measurements, thus improving the resulting ranging estimates. Only a few works have focused on adaptive power control with the intent of improving the reliability or efficiency of communications. Han et al. [25] adapt the TX power to minimize the number of muling devices in communication range in the context of sensor data retrieval: this allows to minimize BLE scanning activities and, therefore, energy consumption. Park et al. [9], [26] have proposed AdaptaBLE, an algorithm adjusting the TX power of BLE devices based on PRR information. Mu et al. [10] have proposed ARTPoS, a system that can adjust the TX power of multiple interfaces in the context of multi-radio IoT platforms. This work, however, focuses mainly on Wi-Fi and ZigBee, and on the problem of optimizing the TX power *while* performing a radio selection.

Hence, to the best of our knowledge, no study explores in detail the effectiveness of different TX power control schemes in BLE, which is essential to understand how to maximize a system's reliability and efficiency. Specifically, there is a lack of works investigating the efficacy of utilizing RSSI information to adjust the TX power (as mandated by LEPC [8]), and comparing it quantitatively with the use of PRR information as

[1]The Bluetooth v5.4 standard has recently defined the PAwR (Periodic Advertising with Responses) feature. The latter allows to build one-to-many bidirectional networks in which receivers can transmit response payloads back to the broadcaster without the need to establish a connection.

[2]Note that energy consumption is especially critical for peripherals, which are often battery-powered and have a lower energy budget than central devices.

[3]BLE foresees 37 data (0–36) and three advertisement channels (37–39).

[4]LEPC support is optional. Two devices supporting this feature *must* use it.

[5]These packets also allow to query information from the remote device, such as the currently-employed TX power and acceptable power reduction.

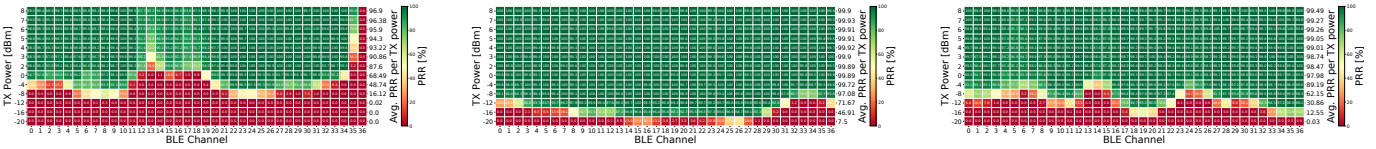

(a) **office-lab** (b) **single-office** (c) **office-office**

Fig. 1: Link characterization of the three device pairs in terms of PRR, as a function of employed BLE channel and TX power.

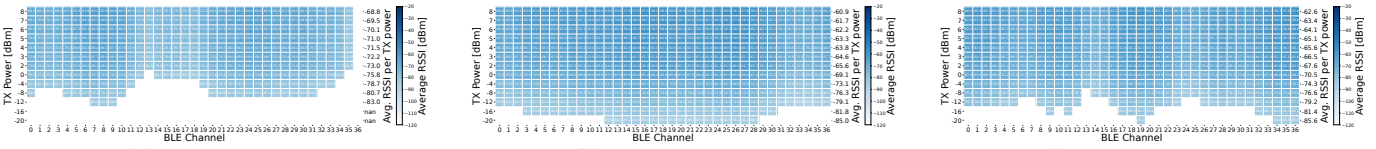

(a) **office-lab** (b) **single-office** (c) **office-office**

Fig. 2: Link characterization of the three device pairs in terms of RSSI, as a function of employed BLE channel and TX power.

proposed in other works [9], [26]. Moreover, existing solutions rely on connection-wide information (*i.e.,* they adjust the TX power equally for all channels). In a previous poster abstract [12], we have argued that *per-channel adaptations* could lead to significant performance improvements, but without backing up this statement with experimental evidence. The next sections of this paper (§ IV–VII) fill this important gap.

Adaptation of low-level BLE settings. A few works have proposed to improve the reliability and energy efficiency of BLE systems by adapting low-level parameters at runtime, such as the CI [27], [28], SL [29], PHY mode [30], and channel map [31], [32]. Jafarizadeh et al. [11] have designed SoftBLE, a software-defined networking framework capable of tuning low-level parameters. In contrast to these works, we focus on adapting the TX power on a per-channel basis.

IV. Experimental Setup and Methodology

To investigate the performance of different TX power control schemes, we collect BLE traffic between several pairs of devices on a public testbed. This real-world data allows us to precisely characterize BLE links and to emulate the effectiveness of different adaptation schemes in the same settings.

Data collection. We run our experiments on the popular D-Cube testbed [33], a public IoT benchmarking facility offering the ability to program tens of nRF52840 DK [34] devices deployed in a university building. We program these devices with Zephyr RTOS v3.0.0 [35] and set up a BLE connection between several pairs of devices (we only operate one pair at a time to avoid interference). Each pair is made up of a central (`C`) and a peripheral (`P`), exchanging data using the 2 Mbps PHY over a connection with CI = 25 ms and SL = 0. In order to characterize the link between `C` and `P`, we let `C` continuously send 207 B-large link-layer packets[6] to `P` every connection event. `C` randomly selects a new TX power (ranging from -20 to +8 dBm) every CE. Tab. I lists the possible TX power values chosen by `C` and their current consumption. To avoid disconnections, `P` uses a fixed TX power of +8 dBm. Each pair of devices exchanges data for 120 minutes, thereby transmitting several hundred thousand packets, during office hours (*i.e.,* in the presence of surrounding Wi-Fi activity).

[6]Each link-layer packet has a payload of 200 B, which contains a sequence number, the TX power employed to send the packet, and padding data.

Device pairs. We run our experiments on pairs of devices deployed at different locations within the D-Cube testbed, in order to diversify and generalize our results. Our evaluation results in § V–VII focus on three device pairs, as summarized in Tab. II. In `office-lab`, `C` and `P` are located in an office and a laboratory, respectively, and are separated by two walls and a corridor. In `single-office`, the two devices are in the same office, in line-of-sight. In `office-office`, `C` and `P` are deployed in two adjacent offices separated by one wall. Fig. 1 and 2 show the characterization of the three device pairs resulting from the data collection in terms of PRR and RSSI as a function of the employed BLE channel and TX power. For example, one can infer from Fig. 1a that, in the `office-lab` scenario, a TX power below -8 dBm and above +2 dBm leads to no reception and to a perfect PRR, respectively, when using channel 22. This information provides an intuitive ground truth of the achievable performance and can be used to investigate the performance of different TX power control schemes.

Evaluation metrics. For each scenario, we compute the average link-layer reliability and power consumption of the radio when using a given adaptation scheme. The reliability is calculated as $PRR(C \rightarrow P) = \frac{\#RX(P)}{\#TX(C)}$, where $\#RX(P)$ is the number of correctly-received link-layer packets at `P` and $\#TX(C)$ is the number of link-layer packets originally sent by `C`. The radio's power consumption ($P_{radio,TX}$) is calculated as the average power used to transmit each individual link-layer packet pkt_{tx}: this allows us to isolate the impact of the employed TX power. Specifically, $P_{radio,TX}(pkt_{tx}) = I_{radio,TX}(pkt_{tx}) \cdot 3.3V$, where $I_{radio,TX}$ is the current consumption of the radio when using a certain TX power value(see Tab. I). Values for $I_{radio,TX}$ are derived using Nordic's *Online Power Profiler for BLE* [36].

TX power [dBm]	-20	-16	-12	-8	-4	0	2	3	4	5	6	7	8
$I_{radio,TX}$ **[mA]**	4.3	4.5	4.6	4.9	5.3	6.3	7.5	8.9	10.2	11.5	12.7	13.9	15.1

TABLE I: TX power values available on the nRF52840 and their respective current consumption ($I_{radio,TX}$).

Scenario	**Distance [m]**	**avg. PRR [%]**		**avg. RSSI [dBm]**	
		+8 dBm	**0 dBm**	**+8 dBm**	**0 dBm**
office-lab	≈ 7	96.9	68.5	-68.8	-75.8
single-office	≈ 5	99.9	99.89	-60.9	-69.1
office-office	≈ 6	99.5	97.98	-62.6	-70.5

TABLE II: Characteristics of the studied device pairs.

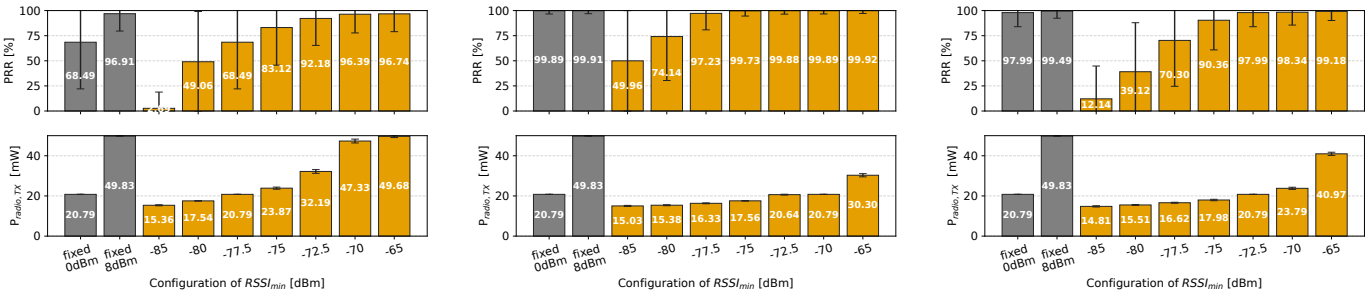

(a) **office-lab** (b) **single-office** (c) **office-office**

Fig. 3: Reliability and energy efficiency of CONN-RSSI in the three test scenarios when using $RSSI_{mrgn} = 0$ dB.

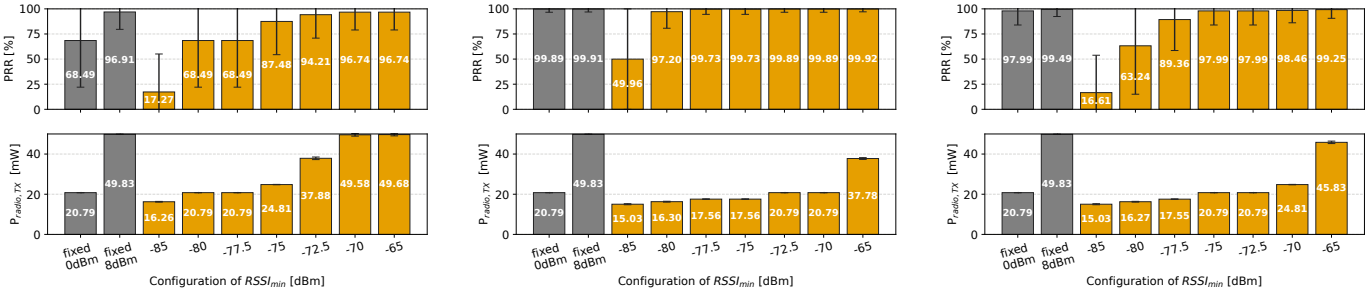

(a) **office-lab** (b) **single-office** (c) **office-office**

Fig. 4: Reliability and energy efficiency of CONN-RSSI in the three test scenarios when using $RSSI_{mrgn} = 2$ dB.

V. RSSI-based TX Power Adaptation

We begin our investigation by analyzing the performance of a TX power control strategy that leverages RSSI information for decision making. Such an algorithm mimics the operation of the LEPC feature, which also needs to adjust the TX power of a device if the RSSI does not fall within a golden range.

Adaptation scheme. Each device (*i.e.,* both C and P) initially selects Zephyr's default TX power (0 dBm) and calculates $RSSI_{curr}$ as a moving average over a pre-defined window size θ. As described in § II, the goal is to keep $RSSI_{curr}$ above a given $RSSI_{min}$ to ensure a good-quality link: as soon as $RSSI_{curr}$ falls below $RSSI_{min}$, the TX power is increased to the next possible value. Once the TX power changes, we refill the moving window before re-calculating $RSSI_{curr}$. As specified by LEPC, to save energy, the TX power can be reduced whenever $RSSI_{curr} > RSSI_{min}$[7]. Since the use of a single threshold (*i.e.,* $RSSI_{min}$) to trigger an increase or decrease of TX power may lead to high fluctuations, we use a hysteresis $RSSI_{mrgn}$. Specifically, we reduce the TX power only if $RSSI_{curr} > (RSSI_{min} + RSSI_{mrgn})$ and increase it when $RSSI_{curr} < RSSI_{min}$. In this work, we study performance when using $RSSI_{mrgn} = 0$ and 2 dB.

We study two variants of this RSSI-based adaptation scheme: either the RSSI moving average is calculated for the entire connection, *i.e.,* using packets received across all channels (CONN-RSSI); or for each individual channel (CHAN-RSSI). That is, we either have a TX power value applied to all channels (CONN-RSSI), or a distinct value per channel (CHAN-RSSI).

[7]Please note that LEPC defines that $RSSI_{curr}$ should be kept below a $RSSI_{max}$ threshold. This is because the use of high TX powers could lead to link failures if devices are in very close proximity [8]. As this is rather unlikely in real-world settings, and as the pairs of devices investigated in this work are several meters apart, we do not account for $RSSI_{max}$.

A. Per-connection adaptation (CONN-RSSI)

In our implementation, we derive $RSSI_{curr}$ using $\theta = 200$: this value was chosen based on empirical analysis and in order not to overwhelm the BLE connection with TX power updates.

The performance of the RSSI-based TX power control for the three scenarios described in § IV is shown in Fig. 3 (for $RSSI_{mrgn} = 0\,dB$) and Fig. 4 (for $RSSI_{mrgn} = 2\,dB$). Specifically, the bar plots show the average and the standard deviation of PRR (top) and $P_{radio,TX}$ (bottom) as a function of $RSSI_{min}$. The latter ranges between -85 dBm (as the nRF52840 defines a valid RSSI range between -90 and -20 dBm) and -65 dBm (as the average RSSI when using the highest TX power never exceeds -60 dBm, see Fig. 2). The first two bars (in grey color)show the baseline performance taken when using a fixed TX power throughout the entire data exchange (*i.e.,* when no adaptation is performed). We choose as baseline the maximum TX power available on the nRF52840 (+8 dBm) and Zephyr's default settings (0 dBm).

As expected, the higher is $RSSI_{min}$, the higher are both PRR and $P_{radio,TX}$. PRR is highest when using the maximum possible TX power at all times (fixed 8 dBm), which however results in the highest energy expenditure[8]. PRR and $P_{radio,TX}$ are identical when using a fixed TX power of 0 dBm and an $RSSI_{min}$ of -77.5 dBm in the office-lab: this makes sense, as the average RSSI when using a fixed TX power of 0 dBm is -75.8 dBm (see Fig. 2a). The same holds true in single-office and office-office: one obtains the same PRR and $P_{radio,TX}$ as when using a fixed TX power of 0 dBm when using an $RSSI_{min}$ of -70, and -72.5 dBm,

[8]Please note that for office-lab the PRR cannot exceed 96.9%, because channel 36 suffers from very poor link quality and no packet is received even when using the highest TX power (see Fig. 1a). In the other two scenarios, instead, the PRR always converges to almost 100%.

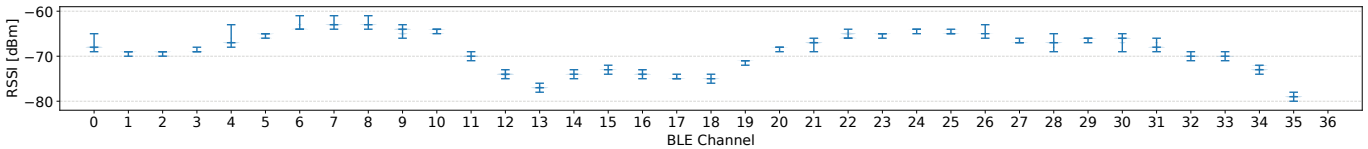

Fig. 5: Distribution of the measured RSSI per BLE data channel in `office-lab` (TX power = +8 dBm). The RSSI in channel 36 is not available, as no packets could be received due to very poor link quality (PRR= 0%).

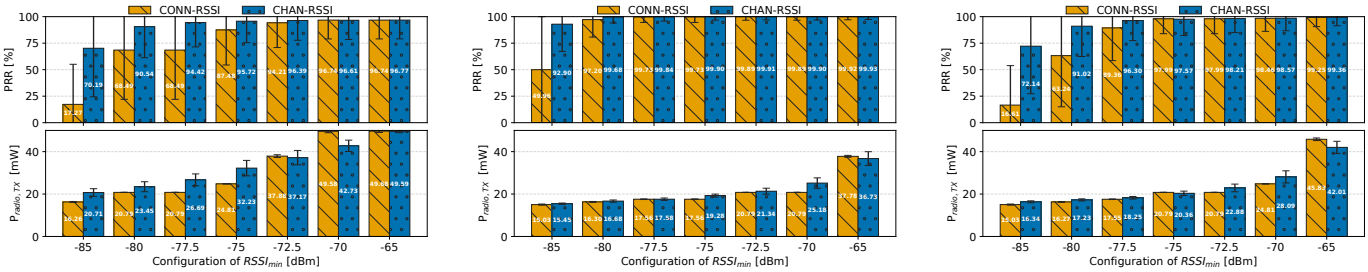

(a) **`office-lab`** (b) **`single-office`** (c) **`office-office`**

Fig. 6: Performance of per-connection (`CONN-RSSI`) and per-channel (`CHAN-RSSI`) adaptation based on RSSI information.

respectively (the average RSSI when using a fixed TX power of 0 dBm is -69.1, and -70.5 dBm, see Figs. 2b and 2c). This, however, also means that it is not possible to select a generic $RSSI_{min}$ that allows to achieve an optimal performance in all scenarios. In fact, the RSSI strongly depends on the environmental conditions and distance between devices; moreover, the RSSI readings of BLE devices are typically not calibrated. Fig. 4c and especially Fig. 4b show that, when devices are close to each other and the link quality is good, a high PRR can be achieved also when using a low $RSSI_{min}$: however, the selection of the RSSI threshold is crucial. For example, in `single-office`, the PRR varies by only 0.25% between $RSSI_{min}$ = -80 and -65 dBm, but $P_{radio,TX}$ grows by 120%. By comparing Figs. 3 and 4 we can see that introducing an hysteresis allows to increase the PRR given the same $RSSI_{min}$.

B. Per-channel adaptation (`CHAN-RSSI`)

We also analyze the performance of RSSI-based TX power control schemes when using per-channel instead of per-connection adaptations, something not considered in earlier works. This makes sense, as the RSSI is frequency-dependent: one can observe variations by up to several tens of dB across different channels. For example, Fig. 5 shows the average RSSI in each BLE data channel in `office-lab`: we can observe that it can vary by up to 15 dB. This hints that averaging the RSSI across all channels may hide precious information to further optimize performance. Essentially, `CHAN-RSSI` enforces a $RSSI_{min}$ for every channel: in `CONN-RSSI`, a channel may exhibit a $RSSI_{curr} < RSSI_{min}$ and still not trigger a TX power change, which may result in a lower PRR. Specifically, in `CHAN-RSSI` we adapt the TX power independently for each channel using $\theta = 5$ (derived to allow a fair comparison with `CONN-RSSI`, *i.e.,* using 200 divided by 37) and $RSSI_{mrgn}$ = 2 dB. Fig. 6 shows a head-to-head comparison of the performance of `CONN-RSSI` and `CHAN-RSSI` in our three test scenarios ($RSSI_{mrgn}$ =2 dB). We can observe that, given a $RSSI_{min}$, the PRR offered by `CHAN-RSSI` is often much superior than that sustained by `CONN-RSSI`, especially for lower thresholds. In Fig. 6a, we can observe that, in the `office-lab` scenario (which is the most challenging of the three from a link quality perspective), to obtain a PRR of 95.72%, `CHAN-RSSI` requires 32.23 mW (and $RSSI_{min}$ = -75dBm), whereas `CONN-RSSI` needs a higher threshold to achieve a PRR of 94.21%, but with a 17% higher $P_{radio,TX}$ (37.88 mW).

VI. PRR-based TX Power Adaptation

We now study the performance of a TX power control strategy that leverages link-layer PRR information for decision making. Such an algorithm mimics the operation of existing approaches in the literature such as AdaptaBLE [9], and follows a similar design rationale as the strategy presented in § V. That is, `C` and `P` initially select Zephyr's default TX power (0 dBm) and calculate the current PRR_{curr} as a moving average over a pre-defined window size θ. If PRR_{curr} is below a threshold PRR_{th}, the TX power is increased to the next available level. If $PRR_{th} \leq PRR_{curr} < (PRR_{th} + PRR_{mrgn})$, the TX power is not changed. For PRR $\geq (PRR_{th} + PRR_{mrgn})$, we reduce the TX power to the next available TX power level. PRR_{mrgn} was empirically set to 5% in our evaluation. As for § V, we study two variants of this adaptation scheme: either the PRR moving average is calculated for the entire connection, *i.e.,* using packets received across all channels (`CONN-PRR`); or for each individual channel (`CHAN-PRR`). We use $\theta = 10$ for `CHAN-PRR` (to have a PRR granularity of at least 10%) and $\theta = 400$ for `CONN-PRR` (to sample $\approx$ 10 samples per channel before updating the TX power).

A. Per-connection adaptation (`CONN-PRR`)

Fig. 7 shows the PRR and $P_{radio,TX}$ of `CONN-PRR` (orange bars) for the three test scenarios and different PRR_{th} configurations. The average PRR sustained is close to the chosen PRR_{th} (or to the maximum achievable PRR by the link when this is lower than PRR_{th}, *e.g.,* for `office-lab`). The

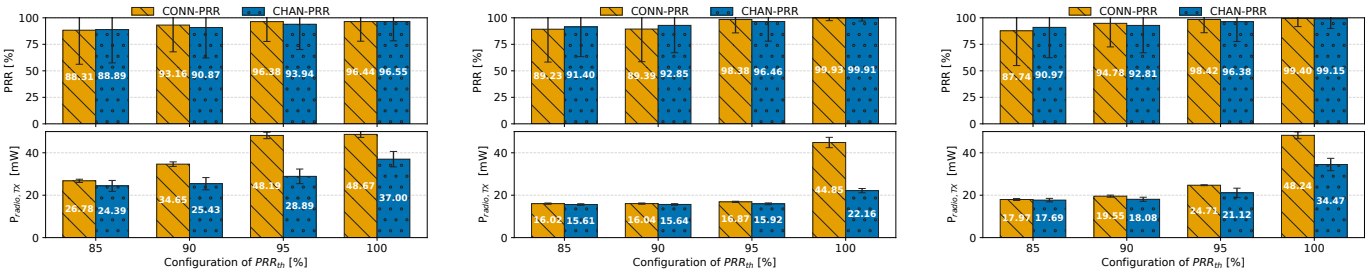

(a) **office-lab** (b) **single-office** (c) **office-office**

Fig. 7: Performance of per-connection (CONN-PRR) and per-channel (CHAN-PRR) adaptation based on PRR information.

power consumption is similar to that observed when running CONN-RSSI; however, we observe that CONN-RSSI requires a lower energy expenditure than CONN-PRR to achieve a comparable average PRR. For example, in office-lab, CONN-PRR achieves a PRR of 96.38% with a $P_{radio,TX}$ of 48.19 mW; CONN-RSSI exhibits a PRR of 96.39%, but with only 37.17 mW (22.8% less power). In single-office, CONN-PRR achieves a PRR of 99.93% with $P_{radio,TX}$ = 44.85 mW; CONN-RSSI exhibits a PRR of 99.92%, but with only 37.78 mW. Similarly, in office-office, CONN-PRR achieves a PRR of 87.74% with c= 17.97 mW; CONN-RSSI exhibits a PRR of 89.36%, but with only 17.55 mW.

B. Per-channel adaptation (CHAN-PRR)

Fig. 7 also shows the PRR and $P_{radio,TX}$ of CHAN-PRR (blue bars) for the three test scenarios and different PRR_{th} configurations. As for RSSI-based approaches, we can clearly observe that *per-channel adaptation* pays off, as it allows to achieve a comparable PRR as CONN-PRR, but with a significantly lower energy expenditure (up to 50.6% lower $P_{radio,TX}$), especially for higher PRR_{th} values. When comparing the energy consumption achieved by CHAN-RSSI and CHAN-PRR in Figs. 6 and 7, we can observe that – as for per-connection adaptation – CHAN-RSSI requires a lower energy expenditure than CHAN-PRR to achieve a comparable PRR.

VII. Putting things together

Next, we distill our findings w.r.t. the performance of RSSI- and PRR-based approaches, as well as to the advantages of per-channel adaptations (§ VII-A). We then propose and implement a new TX power control scheme embedding our observations (§ VII-B), and evaluate its performance (§ VII-C).

A. Distilled findings

The experimental results presented in the previous sections have shown that TX power control schemes based on RSSI and on PRR information are both effective. RSSI-based approaches seem to achieve a comparable reliability to PRR-based schemes, but with a lower energy expenditure. However, they suffer from a crucial drawback: RSSI information is only available if a packet is actually received. Lost packets, which often indicate a poor link quality (and hence the need to increase the TX power to increase reliability), are only accounted for when exploiting PRR information.

Another key advantage of PRR-based approaches is that one can dimension PRR_{th} to match roughly the PRR that should be sustained by the application. In other words, one can fine-tune the reliability of the data exchange such that it meets the requirements of the application at hand: this is something that cannot be done using RSSI-based approaches, as one cannot easily relate the $RSSI_{min}$ threshold to a certain PRR level.

However, a key drawback of PRR-based approaches is that, to know whether reducing the TX power has an effect on reliability, one needs to first experience packet loss: this may not be tolerable in applications requiring short communication delays. Using RSSI-based information (and enforcing a sufficiently high $RSSI_{min}$) allows instead to reduce the TX power while reducing the chances to incur packet loss.

Therefore, we argue that, to harness the best performance, TX power control strategies for BLE should exploit *both* PRR and RSSI information. Moreover, we argue that the TX power should be adjusted per-channel: while this incurs additional memory consumption (one needs to record an array of values for each channel), this is a negligible price when accounting for the significant improvements in energy efficiency compared to the adjustment of TX power on a per-connection basis. However, if a channel is simply poor and reliable communication is not possible, per-channel adaptation alone would end up selecting the highest TX power without sustaining a high PRR. For this reason, informing the adaptation algorithm that such poor channels should be avoided would be beneficial.

B. Proposed control scheme

We hence propose a new TX power control scheme that (i) makes use of *both* RSSI and PRR information, (ii) adapts the TX power on each BLE channel individually, and (iii) updates the used channel map in case of degraded link quality.

Algorithm. We define our control scheme as follows. If (PRR $>$ 99% AND RSSI $>$ -70 dBm), we decrease the TX power ($\delta_{TXP} = -1$). If ($80\% <$ PRR $< 95\%$), we increase the TX power ($\delta_{TXP} = +1$), so to sustain – if possible – a minimum reliability of $\approx$95%. If (PRR $\leq$ 80%), we consider removing a channel from the channel map: thus we specify that if (PRR $\leq$ 70% OR the TX power had just been increased), the channel is removed from the channel map; otherwise we increase the TX power and probe if that helps ($\delta_{TXP} = +1$). In all other cases, we keep the currently-used TX power ($\delta_{TXP} = 0$). The TX power adjustments are made per

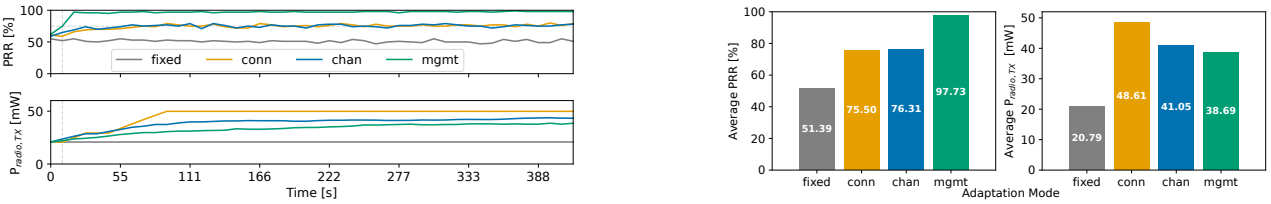

(a) PRR and $P_{radio,TX}$ during the experiment.

(b) Average PRR and $P_{radio,TX}$ for different modes.

Fig. 8: Performance of the proposed *mgmt* TX power control scheme in comparison with three other baselines (`office-lab`).

channel and in a step-wise fashion ($\delta_{TXP} = \{-1, +1\}$)[9]. The proposed scheme hence exploits PRR information to trigger an *increase* in TX power and both PRR and RSSI information to elicit a *decrease* in TX power, as highlighted in § VII-A.

Implementation. We implement the aforementioned scheme on the nRF5840-DK using the Zephyr RTOS. We use a window size θ for PRR and RSSI information of 10 samples per channel[10]. In case per-channel adaptation is not enabled (see § VII-C), we use $\theta = 370$ for comparability (10 samples for each of the 37 BLE data channels = 370 samples). Whenever the number of available channels in the channel map drops below 10, all channels are re-included, as suggested in [31].

In order to make a decision on whether to adjust its TX power, a device needs to receive feedback from its communication partner. For example, in a BLE connection, if the central wants to adjust its TX power, it needs to know what is the link quality perceived by the peripheral, and vice-versa. Feedback can either be in the form of link quality information (*e.g.,* PRR, RSSI), in which case a node can infer the corresponding action autonomously, or in the form of instruction (*e.g.,* increase, decrease, keep the TX power). Such feedback can be provided either on a periodic basis or on-demand. For our proof-of-concept we employ custom link-layer control packets sent at a fixed interval (T_{update} = 370 CI[11]) containing a numeric value (δ_{TXP}) representing the TX power change to be applied.

C. Evaluation

We evaluate the proposed scheme experimentally on the D-Cube testbed in the most challenging scenario (`office-lab`). We use a similar configuration as the one described in § IV, *i.e.,* we set up a pair of nodes exchanging 207-byte packets using CI = 25 ms and the 2 Mbps PHY mode.

We compare the performance of the proposed scheme (*mgmt*) with three baselines: (1) *chan*, which resembles the proposed scheme with per-channel TX power adaptation (same procedure and thresholds), but without updates to the channel map; (2) *conn*, which implements per-connection TX power adaptation and does not update the channel map; (3) *fixed*, which uses a fixed TX power of 0 dBm (Zephyr's default).

We focus on a scenario in which the central (`C`) monitors the link quality and instructs the peripheral (`P`) how to adjust its TX power: this is a common scenario, as central devices have typically abundant energy supply, whereas peripherals are battery-powered. We hence compute the link-layer reliability of communications as $PRR(C \leftrightarrow P) = \frac{\#ACK(P \rightarrow C)}{\#TX(C \rightarrow P)}$, where $\#ACK(P \rightarrow C)$ is the number of acknowledged packets and $TX(C \rightarrow P)$ the number of packets sent by the central. In line with the previous sections, we use $P_{radio,TX}$ to estimate the efficiency of the adaptation. We perform several measurement runs (each 30-minutes long) for each TX power adaptation scheme: we run these sequentially to avoid interference.

Fig. 8b shows the results of our evaluation: the *fixed* approach cannot ensure reliable communication, as a TX power of 0 dBm is insufficient to sustain a high PRR due to the large distance between `C` and `P`. The *conn* and *chan* approach improve the PRR to roughly 75%: however, as they do not filter channels that are unreliable despite the use of high TX power values, it cannot reach a PRR close to 100%. This is, instead, the advantage introduced by the *mgmt* approach, which updates the channel map and sustains a PRR of 97.73%. In line with the previous sections, we can notice that the use of per-channel adaptations (as opposed to per-connection adaptation) allows to decrease the power consumption when sustaining a certain PRR (15.6% less power to obtain a similar PRR between *chan* and *conn*). Fig. 8b shows in the detail the evolution of PRR and $P_{radio,TX}$ over time for the four TX power control schemes, confirming the superiority of *mgmt*. We observe that *mgmt* takes relatively little time to outperform the other solutions: the reason is that it takes only few T_{update} intervals to filter poor channels. The number of blacklisted channels is respectively 0 and 12 after the first two update intervals, and 25 at the end of the tests (blacklisted channels overlap with Wi-Fi channels 1 and 11, which were active in the building where D-Cube is deployed)[12].

In summary, the experiments confirm that our proposed TX power control scheme *mgmt* incorporating all the observations and lessons learned in our experimental campaign (§ IV–VI) is superior to existing solutions (both the one proposed by the BLE standard – LEPC – and the ones in existing literature based only on per-connection PRR adaptation).

[9] This means that we increase or decrease the TX power to the next possible value. In principle, one could use more aggressive strategies adjusting the TX power by several steps, but this goes beyond the scope of this work.

[10] In general, the window size should be longer for short connection intervals and shorter for longer connection intervals to avoid keeping old samples for too long in the buffer. However, when devices are stationary, selecting the window size is less critical than in applications with mobile devices.

[11] When using CI = 25 ms, this corresponds to an update every 9.25 s. We choose 370 to match the window size for the connection-based adaptation.

[12] Although the tests were performed in the same `office-lab` scenario as in § IV–VI, a few months in between have passed (during which we finalized the embedded implementation): results are hence not directly comparable.

VIII. Conclusion and Future Work

In this work, we shed light on the efficacy of different TX power control schemes for BLE. Based on experimental evidence, we distill several recommendations and propose a novel approach that makes use of *both* RSSI and PRR information, adapts the TX power per-channel, and discards poor channels. In the future, we plan to investigate variants of the proposed approach (*e.g.,* with different thresholds) and to replace RSSI with SNR information to prevent calibration-related issues.

Acknowledgements

This work was supported by the TRANSACT project. TRANSACT (https://transact-ecsel.eu/) has received funding from the Electronic Component Systems for European Leadership Joint Undertaking under grant agreement no. 101007260. This joint undertaking receives support from the European Union's Horizon 2020 research and innovation programme and Austria, Belgium, Denmark, Finland, Germany, Poland, Netherlands, Norway, and Spain. TRANSACT is also funded by the Austrian Federal Ministry of Transport, Innovation and Technology under the program "ICT of the Future" (https://iktderzukunft.at/en/).

References

[1] F. Bonavolontá *et al.*, "An Improved Method Based on Bluetooth Low-Energy Fingerprinting for the Implementation of PEPS System," *Sensors*, vol. 22, no. 24, Aug. 2022.

[2] P. Li *et al.*, "BmmW: A DNN-based Joint BLE and mmWave Radar System for Accurate 3D Localization," in *Proceedings of the 19th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*. IEEE, Jun. 2023, pp. 47–54.

[3] Q. Chen and L. Tang, "A Wearable Blood Oxygen Saturation Monitoring System based on Bluetooth Low Energy Technology," *Computer Communications*, vol. 160, no. 1, pp. 101–110, Jul. 2020.

[4] P. C. Ng *et al.*, "COVID-19 and Your Smartphone: BLE-Based Smart Contact Tracing," *IEEE Systems Journal*, vol. 15, no. 4, pp. 5367–5378, Dec. 2021.

[5] M. Giordano *et al.*, "SmartTag: An Ultra Low Power Asset Tracking and Usage Analysis IoT Device with Embedded ML Capabilities," in *Proceedings of the IEEE Sensors Applications Symposium*, Aug. 2021.

[6] Bluetooth SIG, "Bluetooth Core Specification, Revision 5.4," 2023.

[7] M. Spörk *et al.*, "Performance and Trade-offs of the new PHY Modes of BLE 5," in *Proceedings of the International Workshop on Pervasive Systems in the IoT Era (PERSIST-IoT)*. ACM, Jul. 2019, pp. 7–12.

[8] Bluetooth SIG, "Bluetooth Core Specification, Revision 5.2," 2019.

[9] E. Park *et al.*, "AdaptaBLE: Data Rate and Transmission Power Adaptation for Bluetooth Low Energy," in *Proceedings of the IEEE Global Communications Conference (GLOBECOM)*, Dec. 2019, pp. 1–6.

[10] D. Mu *et al.*, "Robust Optimal Selection of Radio Type and Transmission Power for Internet of Things," *ACM Transactions on Sensor Networks*, vol. 15, no. 4, Nov. 2019.

[11] M. Jafarizadeh *et al.*, "SoftBLE: An SDN Framework for BLE-based IoT Networks," in *Proceedings of the 6th International Conference on Internet-of-Things Design and Implementation (IoTDI)*. ACM, Apr. 2021, p. 221–233.

[12] E. Salomon *et al.*, "Improving the Reliability of BLE Communications through Packet-level Adaptations on a Per-Channel Basis," in *Proceedings of the 19th International Conference on Embedded Wireless Systems and Networks (EWSN), poster session*, Oct. 2022, pp. 210–211.

[13] Bluetooth SIG, "Bluetooth Core Specification, Revision 5.0," 2016.

[14] T. Watteyne *et al.*, "Reliability through Frequency Diversity: Why Channel Hopping Makes Sense," in *Proceedings of the 6th Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*. ACM, Oct. 2009, pp. 116–123.

[15] S. Lin *et al.*, "ATPC: Adaptive Transmission Power Control for Wireless Sensor Networks," *ACM Transactions on Sensor Networks*, vol. 12, no. 1, Mar. 2016.

[16] J. Jeong *et al.*, "Empirical Analysis of Transmission Power Control Algorithms for Wireless Sensor Networks," in *Proceedings of the 4th International Conference on Networked Sensing Systems (INSS)*. IEEE, Jun. 2007, pp. 27–34.

[17] Y. Fu *et al.*, "Practical Control of Transmission Power for Wireless Sensor Networks," in *Proceedings of the 20th International Conference on Network Protocols (ICNP)*. IEEE, Oct. 2012, pp. 1–10.

[18] L. H. A. Correia *et al.*, "Transmission Power Control in MAC Protocols for Wireless Sensor Networks," in *Proceedings of the 8th International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*. ACM, Oct. 2005, pp. 282—289.

[19] X. Chen *et al.*, "Saving Energy by Adjusting Transmission Power in Wireless Sensor Networks," in *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, Dec. 2011, pp. 1–5.

[20] M. Bor and U. Roedig, "Lora transmission parameter selection," in *Proceedings of the 13th International Conference on Distributed Computing in Sensor Systems (DCOSS)*. IEEE, Jun. 2017, pp. 27–34.

[21] W. Gao *et al.*, "AdapLoRa: Resource Adaptation for Maximizing Network Lifetime in LoRa Networks," in *Proceedings of the 28th International Conference on Network Protocols (ICNP)*. IEEE, Oct. 2020, pp. 1–11.

[22] U. M. Qureshi *et al.*, "Analysis of Bluetooth Low Energy (BLE) Based Indoor Localization System with Multiple Transmission Power Levels," in *Proceedings of the 27th International Symposium on Industrial Electronics (ISIE)*. IEEE, Jun. 2018, pp. 1302–1307.

[23] P. Barsocchi *et al.*, "Occupancy Detection by Multi-Power Bluetooth Low Energy Beaconing," in *Proceedings of the International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, Sep. 2017.

[24] M.-K. Sie and C.-H. Kuo, "Indoor Location Estimation using BLE Beacon with Multiple Transmission Power Levels," in *Proceedings of the International Conference on Consumer Electronics - Taiwan (ICCE-TW)*. IEEE, Jun. 2017, pp. 323–324.

[25] C.-K. Han *et al.*, "Mobility-Driven BLE Transmit-Power Adaptation for Participatory Data Muling," in *Proceedings of the 24th International Conference on Parallel and Distributed Systems (ICPADS)*. IEEE, Dec. 2018, pp. 962–971.

[26] E. Park *et al.*, "AdaptaBLE: Adaptive Control of Data Rate, Transmission Power, and Connection Interval in Bluetooth Low Energy," *Computer Networks*, vol. 181, Nov. 2020.

[27] M. Spörk *et al.*, "Improving the Timeliness of Bluetooth Low Energy in Noisy RF Environments," in *Proceedings of the 16th International Conference on Embedded Wireless Systems and Networks (EWSN)*, Feb. 2019, pp. 23–34.

[28] T. Lee *et al.*, "CABLE: Connection Interval Adaptation for BLE in Dynamic Wireless Environments," in *Proceedings of the 14th International Conference on Sensing, Communication, and Networking (SECON)*. IEEE, Jun. 2017, pp. 1–9.

[29] M. Spörk *et al.*, "Ensuring End-to-End Dependability Requirements in Cloud-based Bluetooth Low Energy Applications," in *Proceedings of the 18th International Conference on Embedded Wireless Systems and Networks (EWSN)*, Feb. 2021, pp. 55–66.

[30] X. Luo *et al.*, "AptBLE: Adaptive PHY Mode based on K-means Algorithm in Bluetooth5," in *Proceedings of the Wireless Telecommunications Symposium (WTS)*. IEEE, Apr. 2021, pp. 1–6.

[31] M. Spörk *et al.*, "Improving the Reliability of Bluetooth Low Energy Connections," in *Proceedings of the 17th International Conference on Embedded Wireless Systems and Networks (EWSN)*, Feb. 2020.

[32] V. Poirot and O. Landsiedel, "eAFH: Informed Exploration for Adaptive Frequency Hopping in Bluetooth Low Energy," in *Proceedings of the 18th International Conference on Distributed Computing in Sensor Systems (DCOSS)*. IEEE, May 2022, pp. 1–8.

[33] M. Schuß *et al.*, "A Competition to Push the Dependability of Low-Power Wireless Protocols to the Edge," in *Proceedings of the 14th International Conference on Embedded Wireless Systems and Networks (EWSN)*, Feb. 2017, pp. 54–65.

[34] Nordic Semiconductor, "nRF52840 Development Kit, User Guide v1.3," Feb. 2019, https://infocenter.nordicsemi.com/pdf/nRF52840_DK_User_Guide_v1.3.pdf.

[35] The Linux Foundation, "The Zephyr Real-Time Operating System," https://www.zephyrproject.org/.

[36] Nordic Semiconductor, "Online Power Profiler for Bluetooth LE," https://tinyurl.com/p4x9jbyc.