

Power-Domain Interference Graph Estimation for Multi-hop BLE Networks

HAIFENG JIA, UM-SJTU Joint Institute, Shanghai Jiao Tong University, Shanghai, China

YICHEN WEI, UM-SJTU Joint Institute, Shanghai Jiao Tong University, Shanghai, China

YIBO PI, UM-SJTU Joint Institute, Shanghai Jiao Tong University, Shanghai, China

CAILIAN CHEN, Department of Automation, Shanghai Jiao Tong University, Shanghai, China

Traditional wisdom for network management allocates network resources separately for the measurement and communication tasks. Heavy measurement tasks may compete limited resources with communication tasks and significantly degrade overall network performance. It is therefore challenging for the interference graph, deemed as incurring heavy measurement overhead, to be used in practice in wireless networks. To address this challenge in wireless sensor networks, our core insight is to use power as a new dimension for interference graph estimation (IGE) such that IGE can be done simultaneously with the communication tasks using the same frequency-time resources. We propose to marry power-domain IGE with concurrent flooding to achieve simultaneous measurement and communication in BLE networks, where the power linearity prerequisite for power-domain IGE holds naturally true in concurrent flooding. With extensive experiments, we conclude the necessary conditions for the power linearity to hold and analyze several non-linearity issues of power related to hardware imperfections. We design and implement network protocols and power control algorithms for IGE in multi-hop BLE networks and conduct experiments to show that the marriage is mutually beneficial for both IGE and concurrent flooding. Furthermore, we demonstrate the potential of IGE in improving channel map convergence and convergecast in BLE networks.

CCS Concepts: • **Networks** → **Network measurement**; *Network protocol design*; • **Computer systems organization** → **Sensor networks**;

Additional Key Words and Phrases: BLE, interference graph, concurrent transmission, flooding, power control

ACM Reference Format:

Haifeng Jia, Yichen Wei, Yibo Pi, and Cailian Chen. 2024. Power-Domain Interference Graph Estimation for Multi-hop BLE Networks. *ACM Trans. Sensor Netw.* 20, 6, Article 115 (October 2024), 28 pages. <https://doi.org/10.1145/3689635>

The article is an extension of our conference paper accepted at the 2023 International Conference on Embedded Wireless Systems and Networks (EWSN) [19].

This work was supported by the National Natural Science Foundation of China under grant 62201346.

Authors' Contact Information: Haifeng Jia, UM-SJTU Joint Institute, Shanghai Jiao Tong University, Shanghai, China; e-mail: ethanjia@sjtu.edu.cn; Yichen Wei, UM-SJTU Joint Institute, Shanghai Jiao Tong University, Shanghai, China; email: ethepherein@sjtu.edu.cn; Yibo Pi (Corresponding author), UM-SJTU Joint Institute, Shanghai Jiao Tong University, Shanghai, China; e-mail: yibo.pi@sjtu.edu.cn; Cailian Chen, Department of Automation, Shanghai Jiao Tong University, Shanghai, China; e-mail: cailianchen@sjtu.edu.cn.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 1550-4859/2024/10-ART115

<https://doi.org/10.1145/3689635>

1 Introduction

The interference graph, depicting both the communication and interference channel conditions between network nodes, is a key element for resource management in wireless networks. Given an interference graph, we can estimate the interference of each node to other nodes in the network and allocate network resources (e.g., power, time, and frequency) in a proper manner. Extensive resource allocation schemes have been proposed for wireless sensor networks assuming the availability of the interference graphs [7, 15, 22]. However, compared to its usage, the advances in the efficient measurement of interference graphs are still lacking, which greatly hinders the practical use of existing resource allocation schemes in wireless sensor networks. Efficient **interference graph estimation (IGE)** has great potential in unleashing the power of resource scheduling in wireless sensor networks.

Traditional wisdom for network management is to allocate network resources (e.g., time and frequency) separately for the measurement and data transmission tasks. In other words, the measurement and data transmission tasks have to compete for the time and frequency resources. As a result, on the same channel, the measurement tasks, if not done efficiently, will significantly reduce the opportunities for data transmission. Unfortunately, under the traditional wisdom for network management, IGE inherently incurs high measurement overhead: an N -node network, including $O(N^2)$ total links, consumes at least N slots for measurement, where one node sends a reference signal to all the other nodes for channel estimation in each slot [3, 40]. Such high measurement overhead makes the traditional wisdom not scalable for large networks, not to mention network dynamics requiring frequent updates on the interference graph. To reduce the measurement overhead, another popular approach is to characterize interference by modeling based on the propagation environment, where model parameters are pre-determined and no measurement is needed [20]. Similarly, several recent works have assumed the existence of a direct mapping between node attributes, such as geolocation, and resource scheduling decisions, which can be learned from a large amount of network layouts by deep learning [9, 21]. Without active measurements, these approaches lack the ability to track network dynamics.

In this article, we propose to integrate measurement tasks into data transmission tasks, contrary to the traditional wisdom of separating them, such that the measurement tasks can be conducted together with the data transmission tasks, without reserving time resources solely for conducting the measurement tasks. To achieve this, our core insight is to decompose the superposition of the **received signal strengths (RSS)** from multiple concurrent senders into the signal strengths of individually attenuated transmit powers in the data transmission tasks. Specifically, our approach assumes *the linearity of received power*: the received power of a listener is a linear combination of the channel gains and the transmit powers of the senders. By varying the transmit powers of the senders, we can obtain different received powers of the listener. Then, the channel gains can be obtained by solving a group of equations, if the received powers of the listener and the transmit powers of the senders are all known. Meanwhile, the choice of the preceding transmit powers should also ensure that the data transmission tasks not be affected due to carrying the measurement tasks.

Although the core insight is intriguing, it is difficult to find a suitable data transmission task for integration, which should satisfy the following three requirements: (1) it needs to be conducted in a well-synchronized network such that the signal strength measured by the listener in a period can be matched to the set of the senders transmitting in the same period; (2) varying the transmit powers of the senders does not affect, if not improve, the performance of the data transmission tasks; and (3) the received powers from concurrent senders are additive for the linearity assumption of received power to be true. Of the three requirements, the last one is the most challenging to satisfy, requiring not only strictly synchronized concurrent senders for data transmission but

also careful analysis of the hardware features and imperfections of the **commercial off-the-shelf (COTS)** devices.

Fortunately, we find a perfect match between IGE and concurrent flooding, a technique with increasing popularity in the lower-power wireless networking based on **concurrent transmission (CT)**. Riding on the tide of concurrent flooding, IGE can be easily introduced into wireless sensor networks. In concurrent flooding, senders rebroadcast received packets in strictly synchronized time slots, with sub-microsecond synchronization accuracy. This feature of concurrent flooding is a key enabler for the linearity of received power. One side effect of CT is that the **carrier frequency offsets (CFOs)** between senders may result in strong destructive interference corrupting the received packets, if the individually received powers from the senders are close. This issue can be mitigated with fine-grained power control aided by the interference graph. In other words, concurrent flooding carries the task of IGE and, in return, benefits from better interference management by knowing the interference graph. The marriage of IGE and concurrent flooding is mutually beneficial. More importantly, the measured interference graph has the potential to be useful in other network activities, such as scheduling. To realize the preceding benefits, we first verify the assumption of power linearity on the COTS devices (Section 3.3) and analyze the non-linearity issues (Section 3.5). We propose the sufficient condition for IGE and conduct controlled validation experiments in Section 3.7 and Section 3.8. We design network protocols to support simultaneous flooding and IGE (Section 4.3) and power control algorithms to optimize the transmit powers of nodes based on the estimated interference graph (Section 4.2). Extensive real-world experiments are conducted to demonstrate the mutual benefits of marrying IGE and concurrent flooding (Section 5). IGE is beneficial not only for concurrent flooding where all nodes receive and rebroadcast messages using the same channels but also for other network activities in BLE multi-hop networks. To demonstrate this, we evaluate how IGE can be leveraged to accelerate channel map convergence and improve the reliability of convergecast (Section 6).

In summary, we make the following key contributions:

- (1) We propose to marry the challenging measurement task of IGE with the data transmission task of concurrent flooding, to conduct interference estimation simultaneously with concurrent flooding and to improve the performance of concurrent flooding with the measured interference graph.
- (2) We validate the power linearity assumption on the COTS devices under BLE 5 PHYs both experimentally and theoretically. We identify, analyze, and mitigate several non-linearity issues of received power related to hardware and time synchronization.
- (3) We design network protocols to integrate IGE and flooding and optimize concurrent flooding by avoiding strong destructive interference with a power control scheme based on the interference graph. Our approach is implemented atop and compared with BlueFlood, a recent concurrent flooding protocol over Bluetooth.
- (4) We demonstrate two use cases of IGE for channel map convergence and convergecast in BLE networks to show its potentials.

2 Background and Related Work

Interference Graph Estimation. Traditional wisdom for network management allocates separate network resources for the measurement and data transmission tasks, resulting in the competing relation between the two types of tasks. To control the measurement overhead, cellular networks leverage their capabilities in fine-grained resource allocation and only insert the CSI-IM reference signals on a small portion of subcarriers to measure interference [10, 33]. However, low-power COTS devices typically do not have such capabilities and use the whole bandwidth either for

measurement, control, or data transmission [28, 38]. This makes the heavy measurement tasks, such as IGE, not suitable for wireless sensor networks, despite that extensive works have shown the benefits of knowing the interference graph in wireless sensor networks [6, 11]. Our work proposes to avoid the heavy measurement overhead of IGE by integrating it into a popular flooding technique in recent years.

Concurrent Flooding. Glossy achieves fast and efficient network flooding in a strictly synchronized network under IEEE 802.15.4, where senders rebroadcast received packets in each slot [13]. A recent work, BlueFlood, extends concurrent flooding to Bluetooth 5 [1]. Due to the CFOs between concurrent senders, the receiver will see beating patterns with periodic hills (constructive interference) and valleys (destructive interference) [23], which are found under both Bluetooth PHYs and IEEE 802.15.4 [18]. Beating patterns are strong when the RSS from two concurrent senders are close, causing burst errors due to deep valleys with weak signal strengths [4]. When the RSS from one sender is much larger than the rest, the beating patterns are weak and the capture effect dominates, similar to data transmission with no beating [16]. Our work extends the use of concurrent flooding for IGE and, further, avoids strong destructive interference in concurrent flooding with power control based on the interference graph.

Power Linearity. Power linearity is a crucial assumption behind the well-known additive interference model [17] but was first found inaccurate for CT in the work of Son et al. [35]. Subsequent works [25, 26] reported contradictory experimental results and supported the additive signal strength assumption for CT in IEEE 802.15.4. However, prior works are focused on verifying the assumption, without delving into the causes for non-linearity found in their experiments. We extend the experimental study of power linearity to CT in BLE networks and discuss possible causes for non-linearity issues.

Interference-Aware Scheduling. Interference graphs are extensively used for the resource management in wireless networks [5, 39]. An efficient approach to IGE could greatly unleash the power of resource scheduling for wireless sensor networks. For example, with an interference graph, each sender can choose a proper transmit power to control its interference to others [14], and more senders can be scheduled to transmit at the same time for better spatial reuse [8]. Moreover, CT can avoid strong destructive interference with power control. Distributed MIMO allows the transmitters to adjust their signal phases such that the receiving signals are added constructively. However, it requires strict frequency and time synchronization and involves heavy measurement and communication overheads for BLE nodes to act as distributed MIMO [2]. In contrast, our approach conducts IGE in the power domain, which consumes no extra frequency-time resources and is robust to time synchronization error and CFO.

3 Power-Domain IGE

3.1 Overview

Our power-domain IGE builds on the assumption of power linearity—that is, the individual receive powers from multiple transmitters are additive. We undertake a comprehensive experimental study of power linearity on BLE devices (Figure 1). Specifically, we conduct extensive controlled experiments to verify power linearity under various network settings, where the tuning parameters include time synchronization error, transmit power, receive power, and transmission mode. These experiments enable us to (1) identify network settings leading to non-linearity issues and (2) observe that time synchronization and identical transmit data are the necessary but not sufficient conditions for power linearity. To better understand the sufficient conditions for power linearity, we theoretically model CT in BLE networks to gain deeper insights, where factors including CFO, initial phase offsets, and **received signal strength indicator (RSSI)** sampling are considered.




Fig. 1. Overview of our experimental study of power-domain IGE.




Fig. 2. An example of IGE with power control.

Combining those network settings with non-linearity issues and theoretical modeling, we identify four issues and propose mitigation measures. The resulting power-domain IGE is evaluated in real-world experiments.

3.2 Core Insight

An Example. In Figure 2, nodes 1 and 2 send a packet to node 3 concurrently in two time slots. In slot 1, nodes 1 and 2 send to node 3 with the transmit powers of 1mW and 2mW, respectively, and node 3 receives at 3μW. In slot 2, nodes 1 and 2 change their transmit powers to 1mW and 1mW, respectively, and node 3 receives at 2μW. Assuming that the received power of node 3 is a linear combination of the transmit powers of nodes 1 and 2, we can have two equations for the two slots and obtain $h_{13} = h_{23} = 0.001$, where h_{ij} is the channel gain from node i to j .

Linearity Assumption of Received Power. This example makes an important assumption about the linearity of received power, making it possible to estimate the channel gains by solving a system of equations. The linearity of received power is a common assumption for interference modeling in wireless networks [12, 29], but past measurement studies have reported different observations [27, 36]. We will conduct a comprehensive measurement study of the power linearity and present several non-linearity issues in Section 3.5.

Full-Rank Constraint. This example has a unique solution for the channel gains because the two equations are linearly independent. To generalize this, we want the transmit power matrix, with each row being the transmit powers of nodes in a slot, to be full rank, such that the channel gains have a unique solution.

3.3 Power Linearity on the COTS Devices

Two Properties of Linearity. Under the linearity assumption, the received power of node i from two concurrent senders can be expressed as

$$p_i^{rx} = h_{1i}p_1^{tx} + h_{2i}p_2^{tx}, \quad (1)$$

where p_i^{rx} and p_i^{tx} are the received and transmit powers of node i , respectively. This assumption entails two properties of received power, both of which need to be examined on the COTS devices: (1) *proportionality*, which requires the received power to be proportional to the transmit power (i.e., $p_{i \rightarrow j}^{rx} = h_{ij}p_i^{tx}$, where $p_{i \rightarrow j}^{rx}$ denotes the received power of node j solely from node i), and (2) *additivity*, which requires the received powers to be additive (i.e., $p_i^{rx} = \sum_j p_{j \rightarrow i}^{rx}$).




Fig. 3. Experimental setup.




Fig. 4. Proportionality between Tx and Rx powers.

Experimental Setup. We want to examine the two properties of linearity on the COTS devices. To avoid external interference, we conduct our experiments by connecting all nodes with cables. Figure 3 shows the setup with multiple concurrent senders and one receiver, where the attenuators are used to control the received power from each sender at a granularity of 1dB, and the RF combiner is to mix signals from the concurrent senders. Depending on the specific experiments, different numbers of concurrent senders may be used. All nodes are equipped with the Nordic nRF52 series SoCs to study Bluetooth 5. Each node sends a packet of 255 bytes to allow sufficient time for the receiver to take RSSI samples. With a sampling rate of 10^6 samples per second, the receiver collects 1,200 RSSI samples continuously and calculates the average as the received power.

Proportionality between Tx and Rx Powers. We expect the received power to be proportional to the transmit power, whereas hardware imperfections may affect the proportionality. To examine this, we need one sender and one receiver for experiments. Figure 4(a) shows the received powers under different transmit powers and attenuations for nRF52840 SoCs, where the cable and RF combiner together contribute about 3dB attenuation. We can see that the received power increases by almost the same amount in decibels as the transmit power does for all attenuations in the *linear region*, where the transmit power is below or equal to 0dBm and the received power is between -90dBm and -20dBm. The linear region closely matches the valid range (-90dBm to -20dBm) of the received power for the nRF52 series, but non-linearity starts to become significant when the received power is below -90dBm. We are surprised to find that the transmit powers greater than 0dBm severely deviate from the linear relation with the corresponding received powers. Since there is a sudden jump in deviation for the transmit powers greater than 0dBm, the culprit is likely to be the inaccuracy of the transmit power. We repeated this experiment with six additional pairs of nRF52840 SoCs and found that they all had the same issue.

In the linear region, the transmit power has a strong proportional relation with the received power. We calculate the deviations between the actual and expected received powers, where the expected received power is inferred from the transmit power based on the linear relation.




Fig. 5. Power ratios under impacting factors.

Figure 4(b) shows the distribution of power deviations in the linear region, where the percentages of power deviations below and above 0dB are almost equal. There are about 12% of power deviations less than -1 dB, but all power deviations lie between -2 dB and 2 dB. This is consistent with the specifications of the nRF52 series, where the received and transmit powers have an accuracy of ± 2 dB and ± 1 dB in room temperature, respectively [31, 32].

Additivity of Received Power. The additivity assumption of received power is commonly used to model interference in wireless networks [12, 29], but it may not be true on the COTS devices. We want to start by checking this assumption under the simple case of two concurrent senders and one receiver. We are concerned with three possible impacting factors: (1) data similarity (i.e., if the concurrent senders transmit the same data); (2) power delta, the difference in received power between the concurrent senders; and (3) time synchronization (i.e., if the concurrent senders are synchronized). The metric to measure power linearity, referred to as the *power ratio*, is defined as the ratio of the actual received power to the sum of individually attenuated powers from the senders. A power ratio close to 1 indicates strong additivity, and the additivity weakens when the power ratio deviates from 1.

To synchronize the two senders, we ask the receiver to first broadcast a short packet to the two senders and then switch to receive mode. After receiving the short packet, the two senders switch to transmit mode and transmit packets concurrently with specified powers to the receiver. In the unsynchronized case, we apply a random delay to the transmission time of one of the senders, where the delay is chosen to be greater than a symbol time to distinguish from the synchronized case with errors. We conduct experiments with transmit and received powers in the linear region to exclude the preceding disproportionality issues. For nRF52 series SoCs, there are only seven choices for the transmit powers below or equal to 0dBm. We create granular power deltas between the received powers with adjustable attenuators. In cases where senders need to transmit different packets, we simply ask them to transmit all ones and zeros, respectively.

Figure 5 shows the distributions of the power ratios under different power deltas, where CTs under each power delta are repeated 1,000 times with different received powers and with the BLE 1M mode. The bars in the figure show the medians of power ratios and the corresponding error bars show the 5th and 95th percentiles of power ratios. We can see that sending unsynchronized same packets or well-synchronized different packets causes all of the power ratios to be far below 1 regardless of the power delta, indicating weak additivity. In contrast, sending well-synchronized same packets results in a wide range of power ratios with the median equal to 1 under all power deltas. This implies that time synchronization and data similarity are the necessary but not sufficient conditions for the additivity of received power. In other words, if the packets from concurrent senders are different or if the time is unsynchronized, the superimposed signals have weaker power than the power sum of individually attenuated signals, but sending well-synchronized same packets does not guarantee strong additivity. It can also be seen from Figure 5 that the range of power ratios shrinks as the power delta increases, and that the power ratios gradually approach




Fig. 6. Conditions for the additivity of received powers.




Fig. 7. Power ratios under different modes.

to 1 for sending unsynchronized same packets or well-synchronized different packets. However, this does not mean that the additivity of received power increases with the power delta. Instead, when the power delta is large, the additive power is dominated by the larger received power, which diminishes the impact of non-linearity caused by superposition of the smaller one.

To understand the wide range of power ratios when sending well-synchronized same packets, we first look at the CDF of power ratios, as shown in Figure 6(a), where about 12% of power ratios fall beyond the range between 0.9 and 1.1. We will analyze the non-linearity issues in detail later in Section 3.5 and discuss how to identify and remove these issues to reduce the range of power ratios in Section 3.6. Since the operating range of the received power is critical to power linearity as discussed before, we next relate power ratios with the received power. Figure 6(b) and (c) show the power ratios under the strong (-40dBm to -20dBm) and weak (-80dBm to -60dBm) received powers, respectively, where the received powers are divided into equal-length bins of 4dB and the centers of the bins are used as labels for the x - and y -axes. We can see that the power ratio highly depends on the strength of received powers in the region of strong received powers, where the average power ratio gradually decreases from 1.11 to 0.71 when the received powers both increase from -40dBm to -20dBm along the diagonal. The additivity of received powers is the best when both the received powers are centered around -32dBm or when the received powers differ significantly from each other. The power ratio gradually approaches to 1 as the power delta between received powers increases, because the larger received power dominates, which is consistent with the results in Figure 5. For weak received powers, except for the region near (-68dBm , -68dBm), all other regions have power ratios close to 1. In general, apart from the strong received powers greater than -24dBm , weak and strong received powers demonstrate similar additivity. The causes for the high power ratios in the regions of strong and weak received powers are different, which will be discussed later in Section 3.5.

We repeat the preceding experiments under different transmission modes, including the BLE 2M, 500K, and 125K modes, and find similar distributions of power ratios as in the BLE 1M mode. Figure 7 shows that power ratios have a large range including one under all transmission modes

and power deltas. This further confirms that sending well-synchronized same packets are necessary but not sufficient to achieve the additivity of received powers. To achieve the additivity of received powers for Bluetooth, the received powers need to be well controlled. Root cause analysis on the non-linearity issues will be conducted in Section 3.5. By mitigating the non-linearity issues, the additivity of received powers can be further improved, especially for those with small power deltas and thus having strong influence to each other. Figure 7 also shows the distribution of power ratios for 802.15.4, where the average power ratios are close to 1, confirming that power linearity is not limited to Bluetooth.

3.4 Theoretical Explanations for Power Linearity

From the preceding experiments, we observe that data similarity and time synchronization are the necessary but not sufficient conditions for power linearity. We want to explain this by modeling the CT in FSK (frequency-shift keying) systems. Let $s_k(t)$ be the modulated transmit waveform from sender k , which can be expressed as

$$s_k(t) = \text{Re}(b_k(t)e^{j2\pi f_{tx,k}t}),$$

where $f_{tx,k}$ is the carrier frequency of the transmitter and $b_k(t) = a_k(t)e^{j\phi_k(t)}$ is the modulator signal. As Bluetooth employs GFSK (Gaussian frequency-shift keying), $a_k(t)$ is constant and $\phi_k(t)$ can be written as

$$\phi_k(t) = \pi m \int_0^t \sum_{i=0}^{\infty} x_{k,i} g(z - iT_s) dz, \quad (2)$$

where m is the modulation index equal to 0.5 for Bluetooth, $x_{k,i}$ is the i -th bit transmitted by sender k , and $g(t)$ is the Gaussian filter for pulse shaping. If we denote the channel coefficient from sender k to the receiver as h_k , the received baseband signal from N senders can be expressed as

$$r(t) = \sum_{k=1}^N h_k b_k(t - \tau_k) e^{j2\pi \Delta f_k t + \Delta \phi_k^0}, \quad (3)$$

where $\Delta \phi_k^0 = \phi_{tx,k}^0 - \phi_{rx}^0$ is the initial phase offset between sender k and the receiver, $\Delta f_k = f_{tx,k} - f_{rx}$ is the CFO between sender k and the receiver, and τ_k is the time delay from sender k to the receiver.

When sending well-synchronized same packets, we can approximately have $\tau_k = \tau$ and $b_k(t - \tau_k) = b(t - \tau)$ for all k 's. Assuming that senders transmit at the same power (i.e., $a_k(t) = a$), the instantaneous power of received signals can be calculated as

$$|r(t)|^2 = \sum_{k=1}^N a^2 |h_k|^2 + 2 \sum_{k=1}^N \sum_{l=k+1}^N a^2 |h_k| |h_l| \cos(2\pi \Delta f_{k,l} t + \Delta \phi_{k,l}^h + \Delta \phi_{k,l}^0), \quad (4)$$

where $\Delta f_{k,l} = f_{tx,k} - f_{tx,l}$, $\Delta \phi_{k,l}^h = \phi_k^h - \phi_l^h$ is the difference between the phases of h_k and h_l , and $\Delta \phi_{k,l}^0 = \phi_{tx,k}^0 - \phi_{tx,l}^0$ is the difference between the initial phases of senders k and l . We can see that the power of received signals from concurrent senders has a sinusoidal pattern, typically referred to as the *beating pattern*, with a period of $1/\Delta f_{k,l}$, inversely proportional to the difference in carrier frequency between the senders. Equation (4) indicates that the average received power over a multiple of beating cycles is additive, where the integral of the sinusoidal term is zero. However, the average received power on the COTS devices can only be calculated by averaging the RSSI samples. Let T_s be the sampling period, T_b be the beating cycle, and $f(t) = \cos(2\pi \Delta f_{k,l} t + \Delta \phi_{k,l})$,




Fig. 8. Power non-linearity due to ultra-fast beating.

where $\Delta\phi_{k,l} = \Delta\phi_{k,l}^h + \Delta\phi_{k,l}^0$. The sum of RSSI samples can be written as

$$\frac{1}{n_s} \sum_{d=1}^{n_s} |r(T_s d)| T_s = \sum_{k=1}^N a^2 |h_k|^2 + 2 \sum_{k=1}^N \sum_{l=k+1}^N a^2 |h_k| |h_l| \frac{1}{n_s} \sum_{d=1}^{n_s} f(T_s d) T_s, \quad (5)$$

which is the Riemann sum of Equation (4). When $T_s \ll 1/\Delta f_{k,l}$, the Riemann sum well approximates the integral. In our experiments, the RSSI is sampled every $1\mu s$, which is far less than the beating cycles. By taking sufficient samples, we can diminish the impact of the sinusoidal term and approximately have that

$$\frac{1}{n_s} \sum_{d=1}^{n_s} |r(T_s d)| T_s = \sum_{k=1}^N a^2 |h_k|^2. \quad (6)$$

This explains the power linearity seen in our experiments. More discussions on the non-linearity issues will be discussed later in Section 3.5.

When sending well-synchronized different packets, we can still approximately have $\tau_k = \tau$, but the modulator signals, $b_k(t - \tau_k)$'s, are different. In Equation (2), $x_{k,i}$ equals to 1 or -1 to represent a bit. When the sender transmits all ones, $\phi_k(t)$ can be simplified to $\frac{\pi}{2T_s} t$, and similarly, when the sender transmits all zeros, $\phi_k(t) = -\frac{\pi}{2T_s} t$. When two senders transmit all ones and zeros, respectively, we can rewrite Equation (4) by replacing $\Delta f_{k,l}$ with $\Delta f'_{k,l} = f_{tx,k} - f_{tx,l} + \frac{1}{2T_s}$. In our experiments, the typical beating cycles, $1/\Delta f_{k,l}$, are about hundreds of microseconds when sending well-synchronized same packets. When sending well-synchronized different packets, the extra item, $\frac{1}{2T_s}$, in $\Delta f'_{k,l}$ is 500kHz for the BLE uncoded 1M and coded modes, making the beating cycle much smaller than $1/\Delta f_{k,l}$. We believe that the ultra-fast beating cycle is the culprit for the power non-linearity when sending well-synchronized different packets.

To understand the culprit, we examine the beating patterns between sending the same and different packets. We let two concurrent senders send all ones and zeros, respectively, with the same transmit power and mode. Figure 8(a) shows an example of the beating patterns with the BLE 1M mode, where the RSSI sampling period is $1\mu s$. We can see a clear beating pattern when sending the same packets, but no beating pattern can be observed when sending different packets. More importantly, half of the RSSI samples are above $2\mu W$ in the beating pattern, but only 24.4% of samples are above the same threshold when sending different packets. We suspect that this is because the **automatic gain control (AGC)**, used for stabilizing the received power, fails to adapt to the ultra-fast beating. To verify this, we repeat the experiment with the BLE 2M mode, where $\Delta f'_{k,l}$ further increases to about 1MHz. Figure 8(b) shows the average power ratios under different received powers for BLE 1M and 2M. The average power ratios under BLE 2M are always lower

Table 1. List of Non-Linearity Issues

No.	Issues	Details
1	Power saturation	Sum of Rx powers exceeds -20dBm
2	Overshoot	AGC output overshoots desired level
3	RSSI sample granularity	Granularity of the RSSI samples is 1dB
4	Time sync error	Sync error between concurrent senders




Fig. 9. Non-linearity issues breakdown.

than those under BLE 1M. This implies that faster beating results in lower power ratios, which confirms that AGC may over-react to ultra-fast beating and cause lower additive power.

3.5 Non-Linearity Issues Breakdown

We take an experiment-based approach to discover non-linearity issues. As in Section 3.3, we have conducted extensive controlled experiments under various settings of CT and then collected all the network settings with non-linearity issues. For each of network settings, we repeated the experiments and collected microsecond-level time series of RSSI samples for detailed analysis. We discovered four issues in these RSSI sample series by comparing them with the ideal beating patterns, theoretically modeled in Section 3.4. Table 1 lists all the issues and the detailed descriptions.

Issue 1: Received Power Saturation. As mentioned before, the maximum nominal received power for the nRF52 series SoCs is -20dBm , beyond which the linear increase of received power suddenly slows down. This indicates that the power linearity will be impaired if the received power saturates. Suppose that the individually received powers from two concurrent senders are denoted as P_1 and P_2 . From Equation (4), we know that if $P_1 = P_2 = P$, the beating pattern will oscillate between 0 and $4P$, where the maxima is four times the individually received power, P . This implies that superimposing two individually received powers above -26dBm could saturate the receiver. This agrees with our observations in Figure 6(b), where the received powers above -26dBm result in the power ratios to be less than 1. Figure 9(a) shows an example of the received power saturation, where the two individually received powers are both close to -21dBm and the sinusoidal pattern has peaks exceeding the -20dBm limit. We can see that the RSSI samples above $10\mu\text{W}$ (i.e., -20dBm) are suppressed, where the peak is $20\mu\text{W}$, which is very close to the desired average, $2P$, and much less than the desired maxima, $4P$.

Issue 2: Overshoot. AGC aims to stabilize the output power by dynamically adjusting its power gain. However, when it underestimates the strength of received power, an overlarge power gain may be used, causing the output power overshoot. We observed both short- and long-term overshoots in our experiments. The short-term overshoot typically causes a sudden jump to a certain RSSI sample and takes several microseconds to return to the normal level, whereas the long-term overshoot may last for a beating cycle and affect all the RSSI samples included. Short-term overshoots could happen anywhere in the beating patterns but are mostly seen near the peaks in our experiments. In Figure 9(b), short-term overshoots occur at the beginning of the first and last peaks,

and a long-term overshoot lasts for more than half of the beating cycle. When an overshoot occurs, the original power could be amplified by multiple times. Since the received power is calculated by averaging the RSSI samples, overamplified RSSI samples will cause the measured received power to be larger than the actual one. This is the major cause for the power ratios to be greater than 1 when the individually received powers from two concurrent senders are near -68dBm (see Figure 6(c)).

Issue 3: RSSI Sample Granularity. In our experiments, nRF52 series SoCs only allow measuring the RSSI samples at the granularity of 1dB. In most cases, if the received power is rounded up or down to the closest decibel evenly, this should not affect the power ratio. However, when the peak RSSI samples take a significant portion of all samples and are all rounded up, the power ratio may become overlarge. Figure 9(c) shows an example of the large power ratio due to the course-grained RSSI sample granularity. About 16% of samples are at the peaks, which is larger than the desired peak for 15.6%, where the desired peak is calculated based on Equation (5).

Issue 4: Impact of Time Synchronization Errors. In our previous experiments, we synchronize concurrent senders by asking them to switch to transmit mode once receiving the synchronization packet from the receiver. However, due to the hardware imperfections, concurrent senders can hardly be perfectly synchronized and will still experience synchronization errors. In this experiment, we need to accurately measure the time synchronization error to study its impact on power linearity. Specifically, suppose that two senders start to transmit packets at t_1 and t_2 , respectively. The time synchronization error is then $|t_1 - t_2|$, which can be measured by comparing the trigger times of the same ratio-registered event between the senders. We use the ADDRESS event in our experiment, which is immediately triggered by the hardware when the access address of a packet is sent out [31, 32]. To avoid CPU delays, we use the PPI (Programmable Peripheral Interconnect) in nRF52 series SoCs to toggle the GPIO pins once the targeted ratio-registered event is triggered, and calculate the difference between the trigger times on the logical analyzer. We evaluate the impact of time synchronization errors on power linearity with experiments as follows.

We repeat the preceding experiments of sending well-synchronized same packets by introducing a random delay below 1,000 ns to one of the two concurrent senders. Figure 9(d) shows the distribution of power ratios under different time synchronization errors with the BLE 1M mode, where the x-axis labels present the centers of the bins with a width of 200ns. It can be seen that the distribution of power ratios remain relatively stable when the time synchronization error is less than half of the symbol time. However, as the time synchronization error continues to increase, the power ratios begin to decrease. It is interesting that although the senders transmit the same packets, the power linearity still decreases with the synchronization error. This implies that as the synchronization error exceeds half of the symbol time, inter-symbol interference becomes severe. We suspect that this is because of data whitening used in Bluetooth, which is often applied to the information bit sequence to make zeros and ones appear as even as possible. In other words, although our information bit sequence comprises all ones, the resulting bit sequence after data whitening will contain both zeros and ones. When the time synchronization error is about a symbol time, it is equivalent to senders transmitting different packets.

3.6 Achieving Power Linearity

Based on our preceding observations, we present methods to mitigate the non-linearity issues and discuss the power linearity under multiple concurrent senders.

Mitigating Non-Linearity Issues. We want to mitigate the non-linearity issues from three aspects: (1) limiting the operating range of the transmit and received powers to be within the linear region (i.e., the transmit and received powers are no more than 0dBm and -20dBm , respectively), (2) controlling the time synchronization error to be within one-fourth of the symbol time, and (3)




Fig. 10. Power ratios before and after mitigation.




Fig. 11. Impact of dense concurrent senders.

removing abnormal RSSI samples due to overshoot. As shown in Figure 9(b), since we have multiple beating cycles, we can remove abnormal RSSI samples by selecting the best beating cycles. This can be done by fitting the RSSI time series with a cosine function and then selecting the two consecutive beating cycle that fits the best, where two consecutive cycles are used to mitigate the error in estimating the beating period. However, non-linear curve fitting may consume the limited computing resources of the embedded systems, making it impractical for implementation. Considering that abnormal RSSI samples only cause overlapped power ratios greater than 1.2 for 0.3% of CTs in our experiment, we employ the first two methods to achieve power linearity in our system implementation (Section 5). Figure 10 shows the ranges of power ratios before and after the preceding methods are applied. We can see that the most significant improvement is for power ratios with a delta less than 1dB. By removing the outliers, the 5th and 95th percentiles of the power ratios are greatly improved.

Dense Concurrent Senders. Real-world applications may involve multiple concurrent senders. We want to understand if power linearity still holds for the cases of multiple senders. From Equation (4), we know that the power linearity should still hold if the sampling duration is a multiple of the longest beating cycle. We repeat the preceding experiments with more concurrent senders, where the receiver continuously collects RSSI samples for $3,000\mu s$. Figure 11 shows the range of power ratios under different numbers of concurrent senders, where the transmit power is selected randomly from $[0, -4, -8, -12, -16, -20, -40]$ dBm. The median of power ratios remain stable as the number of concurrent senders increases, indicating that power linearity holds for dense concurrent senders. Looking at the beating pattern between each sender and the receiver, we find that their beating cycles are between $200\mu s$ and $400\mu s$, and that the duration of our RSSI samples includes several of their beating cycles.

3.7 Full-Rank Constraint

Suppose that there are n senders in a time-slotted multi-hop network. Each node can measure the RSS in a time slot in which it does not transmit data. Let the transmit and received powers of node




Fig. 12. Feasibility of IGE on the COTS devices.

i at the j -th slot be denoted as $p_i^{tx}[j]$ and $p_i^{rx}[j]$, respectively. The channel gain from node i to node j is h_{ij} , and the channel gain vector, $\mathbf{h}_i = [h_{1i}, \dots, h_{ni}]^T \in \mathbb{R}^n$, includes the channel gains from all senders to node i . At a time slot, if the power linearity holds, we can write the received power as a linear combination of the channel gains and the transmit powers (i.e., $p_i^{rx}[j] = \sum_{z=1}^n h_{zi} p_z^{tx}[j]$). We assume that the channel coherence time is much larger than the slot time. By varying the transmit powers of the senders across time slots, we can obtain multiple received powers for node i . Let $\mathbf{p}_i^{rx} \in \mathbb{R}^m$ be the received powers of node i in m time slots, and let $\mathbf{P} \in \mathbb{R}^{m \times n}$ be the corresponding transmit power matrix of the senders, where each row of the matrix dictates the transmit powers of senders in a time slot. We can compute $\mathbf{h}_i$ by solving a constrained linear least square problem for each node—that is,

$$\min_{\mathbf{h}_i} \|\mathbf{P}\mathbf{h}_i - \mathbf{p}_i^{rx}\|^2, \text{ s.t. } h_{ij}^{\min} \leq h_{ij} \leq h_{ij}^{\max}, \forall i, j. \quad (7)$$

In our experiments, we set $h_{ij}^{\min}$ and $h_{ij}^{\max}$ based on the choice of the transmit powers of nodes such that the received powers fall within the operating range of the Nordic nRF SoCs (see Section 5.1). If the rank of the transmit power matrix, $\text{rank}(\mathbf{P})$, is equal to n (i.e., $\mathbf{P}$ is *full rank*), we can obtain a unique solution for $\mathbf{h}_i$. For $\text{rank}(\mathbf{P})$ to be greater than n , the number of time slots, m , should be greater than or equal to n .

3.8 IGE: Controlled Experimental Study

Before discussing how to achieve IGE via CT-based flooding, we first want to validate with experiments that IGE is feasible on the COTS devices. To this end, we use five senders and one receiver in our setup, where the senders are synchronized by a short synchronization packet from the receiver and transmit packets with the BLE 1M mode, as before. To construct a full-rank transmit power matrix, we ask each sender to pick a random transmit power in the linear region for each transmission after the time is synchronized. This process is repeated until we have sufficient measurements for each combination of transmit powers. Let the i -th combination be $\mathbf{c}_i = [p_1^{tx}[i], \dots, p_n^{tx}[i]]^T$. We need at least n combinations to construct a full-rank matrix for n senders. We calculate the *channel gain error* as $|\hat{h}_{ij} - h_{ij}|$, where $\hat{h}_{ij}$ is the estimated channel gain for h_{ij} in decibels.

With the preceding measurement data, we first want to understand if averaging multiple measurements for the same combination of transmit powers helps the channel gain estimation. Figure 12(a) shows the channel gain errors when different sample sizes are used to calculate the average transmit powers for each combination. There is no significant improvement for using a larger sample size, which implies that our measured received powers for the same node are stable. We continue this experiment using single samples of received powers to estimate channel gains. Based on the perturbation theory for linear systems, we know that the estimation error of channel gains, $\hat{h}_{ij}$, depends on the condition number of $\mathbf{P}$ [24]. We thus estimate the channel gains using different combinations of transmit powers (or *transmit power vectors*) and obtain a relation between the estimation error and the condition number. We can see from Figure 12(b) that the average channel




Fig. 13. Integrating IGE with flooding.

gain error decreases together with the condition number. When the total number of transmit power vectors forming the full-rank matrix is more than 11, the condition number starts to decrease at an extremely slow rate. This agrees with the perturbation theory in that a large condition number amplifies the measurement errors and results in a large estimation error. When choosing among candidate transmit power matrices, we prefer the one with a smaller condition number.

The perturbation also suggests that it is more difficult to estimate smaller channel gains. To verify this, we use the adjustable attenuators to create different combinations of channel gains from the five senders to the receiver. In total, 30 combinations of channel gains are created and each is estimated using a full-rank matrix. Based on Figure 12(b), we use 11 transmit power vectors for channel gain estimation. Figure 12(c) shows the estimation error for different magnitudes of channel gains, where the channel gains in each combination are ranked from 1 to 5 in descending order. We can see that the largest channel gains experience the least estimation errors, of which the largest two channel gains can be estimated with the median error about 1dB, the 90th percentile error less than 3dB, and the 95th percentile error less than 5dB. If the available transmit powers are more than 3dB apart, the channel gain can then be used to choose a proper transmit power with a high probability. Figure 12(d) shows the distribution of the estimation errors for Bluetooth. Sending different or unsynchronized data results in much worse estimation errors.

4 Integration of IGE and Concurrent Flooding

This section presents how to achieve the mutual benefits for both IGE and concurrent flooding from their integration such that (1) the interference graph can be estimated simultaneously with flooding using the same frequency-time resources and that (2) the performance of concurrent flooding can be improved with the estimated interference graph.

4.1 Integrating IGE with Flooding: Overview

We focus on a time-slotted multi-hop network, where an initiator disseminates information to the network via concurrent flooding. There are two major management tasks in the network: IGE and network performance optimization. The IGE is conducted in the power domain to estimate the interference graph, and the estimated interference graph is then used to optimize the transmit powers of nodes for better network performance.

Figure 13 shows an example of periodic flooding round after round, each including multiple time slots. The IGE process consists of two subprocesses: measurement and reporting. The measurement process spans several rounds of flooding, each starting with the initiator flooding a packet to the entire network, whereas the reporting process requires each node to report their local measurements back to the initiator. After collecting local measurements from all other nodes, the initiator updates the estimated interference graph and runs the power control algorithm to generate a new power allocation plan (Section 4.2), aiming to optimize the performance of flooding.¹ Then, the new plan is disseminated throughout the network with flooding. Once receiving the plan, each

¹The estimated interference graph can be used for other network activities beyond flooding.

node updates its transmit power for flooding and retains the transmit power until the next IGE process. To adapt to the changing network conditions, the IGE process is conducted periodically, where the *update interval* is the time between two consecutive IGE processes.

In the measurement process, nodes receive and forward packets as in typical flooding in each round of flooding, except that they need to monitor the signal strengths of received powers when not in transmit mode and to adjust transmit powers as required by the power-domain IGE technique for measurement purposes (Section 4.3). In the reporting process, as shown in Figure 13, the local measurements of nodes can be opportunistically piggybacked onto the upstream dataflows toward the initiator (dashed boxes) and only the last reporting event (solid box) within the IGE process is required. If there is no opportunity for piggybacking before the last flooding event, local measurements collected over several rounds will have to be reported all at once to the initiator. Detailed analysis on the measurement overhead will be discussed in Section 4.4.

4.2 Performance Optimization for Flooding: Power Control Algorithm

Given the estimated interference graph from the IGE process, we can coordinate the transmit powers of nodes to optimize the network performance in the flooding-alone period. More specifically, we want to mitigate the strong destructive interference in CT-based flooding, which may corrupt the received packets. It has been shown in the work of Al Nahas et al. [30] that when two concurrent senders transmit the same data, the **packet error rate (PER)** of the receiver is high due to destructive interference when the individually received powers from the senders are close. We thus want to maximize the power deltas of receivers such that there is a dominant transmitter for each receiver. As the flooded packet is propagated through the network hop by hop, we determine the transmit powers of nodes at the i -th hop to maximize the power deltas of the $(i + 1)$ -hop nodes while considering the interference from nodes at the $(i - 1)$ -th and earlier hops. We formulate the power allocation problem as follows. Suppose that a set of senders, $\mathcal{S}$, are transmitting to a set of nodes, $\mathcal{M}$, with a set of interferers, $\mathcal{I}$, whose transmit powers cannot be controlled. Let $\mathcal{P}$ be the set of available transmit powers. The problem can be written as

$$\max_{\{p_i^{tx}\}_{i \in \mathcal{S}}} \Delta \quad (8a)$$

$$\text{s.t. } p_i^{tx} \in \mathcal{P}, \forall i \in \mathcal{S}, \quad (8b)$$

$$\log \left(\frac{\hat{h}_{ij} p_i^{tx}}{\sum_{d \in \mathcal{S} \setminus \{i\}} \hat{h}_{dj} p_d^{tx} + \sum_{d \in \mathcal{I}} \hat{h}_{dj} p_d^{tx}} \right) \geq \Delta, \forall j \in \mathcal{M}, \exists i \in \mathcal{S}, \quad (8c)$$

where Δ is the power delta between the dominant received power from one sender and the sum of those from the rest and $\hat{h}_{ij}$ is the estimated value for h_{ij} . This formulation is a non-linear integer programming problem, aiming to maximize the minimum power delta of the receivers. Equation (8c) ensures that each receiver has at least one dominant transmitter. Since the number of available transmit powers is limited in the COTS devices, we simply solve this problem with a depth-first branch-and-cut approach, where Δ is first initialized to negative infinity. During the depth-first search, Δ is updated if a larger minimum power delta is found on a certain branch, and a branch is cut if it is impossible to find a minimum power delta greater than Δ by searching more nodes on that branch.

For a multi-hop network, we compute the allocated powers for nodes at each hop separately, including the last-hop ones. Optimizing the transmit powers of the last-hop nodes aims to improve the packet reception rate of nodes experiencing packet loss at the second-to-last hop. As the power-domain IGE only estimates the downstream channels from the initiator to the last-hop




Fig. 14. An example of flooding with IGE.

nodes, we assume channel symmetry between the last two hops and reuse the estimated downstream channels to compute the transmit powers of the last-hop nodes. It is noticeable that the power control algorithm needs as input an estimated interference graph. We require the network to conduct point-to-point channel estimation for each link when it starts for the first time. With this estimated interference graph, a full-rank transmit matrix can be computed and, afterward, the interference graph can be updated periodically as in Figure 13.

4.3 Measurement Process: Power-Domain IGE

The power-domain IGE measures the interference graph by adjusting the transmit powers of nodes optimized for flooding. To mitigate the impact of IGE on flooding, we propose the *sequential power adjustment* method to adjust the transmit power of each node in turn. To improve the scalability of the power-domain IGE, we further propose the *cross-hop interference decoupling* method to allow nodes at each hop to adjust transmit powers independently. Figure 14 shows an example of the power-domain IGE in a three-hop network. In each time slot, we use a *communication graph* to describe the relation between senders (solid circles) and receivers (empty circles), where the data transmission and interference links are represented by the solid and dashed arrows, respectively. The sequential power adjustment is conducted across rounds, whereas the cross-hop interference cancellation is conducted across slots in the same round.

4.3.1 Sequential Power Adjustment. This method requires that only one node adjusts its transmit power in each round and that nodes adjust their transmit powers in turn across rounds. As shown in Figure 14, in the first round, node 1 adjusts its transmit power to 2mW, whereas node 2 remains transmitting at 1mW as in previous rounds. In the second round, node 1 switches back to transmit at 1mW, whereas node 2 adjusts its transmit power to 2mW. Combining the transmit powers of nodes at rounds 1 and 2, we can form a full-rank transmit matrix for nodes at the $(i - 1)$ -th hop. To improve the scalability of IGE, the power allocation for nodes at different hops is conducted independently, as will soon be discussed in Section 4.3.2.

Condition for the Full-Rank Transmit Power Matrix. We next discuss the conditions to ensure that the transmit power matrix is full rank under the sequential power adjustment. Specifically, let $\mathbf{P}_b = [\mathbf{p}_1^b, \dots, \mathbf{p}_n^b]$ and $\mathbf{p}_i^b = [p_i^b[1], \dots, p_i^b[m]]^T$ be the transmit power matrix of nodes at the same hop for flooding, where n is the number of nodes, $p_i^b[j]$ is the transmit power of node i at the j -th round, and m is the total number of rounds ($m \geq n$). Since the allocated transmit powers of nodes for flooding remain the same across rounds, we have $p_i^b[j] = p_i^b[k], \forall j \neq k$. Let $\Delta\mathbf{P} = \text{diag}(\Delta p_1[1], \dots, \Delta p_n[n])$ be the transmit power adjustment for IGE, indicating that node i

adjusts its transmit power at the i -th round. The total transmit power matrix is $\mathbf{P} = \mathbf{P}_b + \Delta\mathbf{P}$. We want to carefully choose $\Delta\mathbf{P}$ to ensure that $\mathbf{P}$ is full rank.

THEOREM 4.1. *To ensure that $\mathbf{P}$ is full rank, we need to choose $\Delta p_i[i]$'s such that $\sum_{i=1}^n \frac{p_i^b[1]}{\Delta p_i[i]} \neq 0$.*

PROOF. Let $\mathbf{v} = [p_1^b[i], \dots, p_n^b[i]]^T$ and $\mathbf{u} = [1, \dots, 1]^T$. We can write $\mathbf{P} = \mathbf{u}\mathbf{v}^T$. Based on the Sherman–Morrison formula, if $\mathbf{P}^{-1}$ is invertible and $\mathbf{v}^T \Delta\mathbf{P}^{-1} \mathbf{u} \neq 0$, then $\Delta\mathbf{P} + \mathbf{u}\mathbf{v}^T$ is full rank. In other words, we want to make sure that $\mathbf{v}^T \Delta\mathbf{P}^{-1} \mathbf{u} = \sum_{i=1}^n \frac{p_i^b[i]}{\Delta p_i[i]} \neq 0$. $\square$

Our approach needs at least n rounds of flooding to measure the interference graph for a network of n nodes. This approach easily scales to larger networks using more rounds of flooding for measurement, at the cost of prolonging the IGE process. The transmit power matrix $\mathbf{P}$ is robust to packet loss and may be rank deficient only under heavy packet loss, which will be demonstrated with experiments in Section 5.

Choice of Power Adjustment. Theorem 4.1 provides the condition for the transmit power matrix to be full rank, which ensures the unique solution of channel gains. However, accurate channel gain estimation requires careful selection of $\Delta\mathbf{P}$. A larger power adjustment incurs a larger change in the received power, which can combat the measurement errors due to hardware imperfections and thermal noise. Let $\mathbf{v}$ be the measurement error. We have that

$$\mathbf{P}\mathbf{h}_i = (\mathbf{P}_b + \Delta\mathbf{P})\mathbf{h}_i = \mathbf{p}_i^{rx} + \mathbf{v} \Leftrightarrow \Delta\mathbf{P}\mathbf{h}_i = \Delta\mathbf{p}_i^{rx} + \mathbf{v},$$

where $\Delta\mathbf{p}_i^{rx}$ is the change in received power due to power adjustment. Based on the perturbation theory, we can bound the estimation error of channel gains as

$$\frac{\|\hat{\mathbf{h}}_i - \mathbf{h}_i\|}{\|\mathbf{h}_i\|} \leq \kappa(\Delta\mathbf{P}) \frac{\|\mathbf{v}\|}{\|\Delta\mathbf{p}_i^{rx}\|}, \quad (9)$$

where $\hat{\mathbf{h}}_i$ is the estimated channel gain vector of node i and $\kappa(\Delta\mathbf{P})$ is the condition number of $\Delta\mathbf{P}$. It can be seen that the estimation error of channel gains can be reduced by decreasing the condition number of $\Delta\mathbf{P}$ and increasing the power adjustment. However, increasing the power adjustment causes the adjusted power to deviate further from the optimized transmit power for flooding, resulting in a higher PER during the IGE process. This implies a tradeoff in network performance between the IGE and the flooding-alone periods. As the flooding-alone period is much larger than the IGE process, we prefer large power adjustment for the accurate estimation of channel gains for the flooding-alone period. This tradeoff relation will be demonstrated with experiments in Section 5.1. As the condition number of a matrix is the ratio of its largest and smallest singular values, we can control the condition number of $\Delta\mathbf{P}$ by enforcing all nodes to use the same power adjustment. As implied by Figure 12(b), we can also increase the dimension of transmit power matrix to reduce the condition number, but this also requires using a larger number of rounds for the IGE process.

4.3.2 Cross-Hop Interference Decoupling. In a multi-hop network, the packet is flooded hop by hop across the network. To further shorten the IGE process, we want to decouple the interference between cross-hop nodes such that nodes at each hop only need to care about the signals from their direct upstream nodes. After the cross-hop interference is decoupled, we can adjust the transmit powers of nodes at each hop independently to measure the channel gains from these nodes to their downstream nodes. The total number of rounds needed for IGE is thus reduced to $\max_i \{n_i\}$, where n_i is the number of nodes at the i -th hop.

Figure 14 shows how to decouple cross-hop interference by comparing the communication graphs across different slots in the same round. In slot 1, nodes 1 and 2 at the $(i-1)$ -th hop begin to

rebroadcast a packet at the same time. Nodes 3 and 4 at the i -th hop can receive the packet as they fall within the communication range of nodes 1 and 2, whereas nodes at hop $i + 1$ can only sense the weak signal strength from nodes 1 and 2 for being beyond the communication range. In slot 2, nodes at the i -th hop have received the packet and begin to rebroadcast it, together with nodes 1 and 2. We assume that nodes 1 and 2 use the same transmit power for both slots 1 and 2. Based on the additivity of received power, since hop- $(i + 1)$ nodes have measured the signal strength from nodes 1 and 2 in slot 1, by subtracting the interference from the total received power in slot 2, they can obtain the strength of signals solely from hop- i nodes. This implies that the cross-hop interference can be decoupled by requiring nodes at the same hop to maintain their transmit powers across time slots in the same round. With interference decoupling, nodes at each hop can estimate the channel gains from themselves to their next hops independently.

4.4 Reporting Process: Overhead Analysis

The preceding measurement process requires us to record the communication graph in each time slot. Specifically, we need to know the status of each node, either in listen, transmit, or idle mode, in each time slot. When a round starts, each node except for the initiator stays in listen mode until it receives the flooded packet and then rebroadcasts it for N_{tx} consecutive rounds. This indicates that only the reception time of the flooded packet needs to be recorded. A round consisting of n_s slots would take each node $\lceil \log n_s \rceil$ bits to represent the reception time. Since we focus on flooding, nodes at the i -hop need to take i copies of RSS measurements. Let b be the number of possible RSS values,² each of which can be represented with $\lceil \log_2 b \rceil$ bits. Let d be the number of total hops. The total communication overhead of our approach for measurement data collection can be estimated as

$$O_{our} = \lceil \log n_s \rceil n_r \sum_{i=1}^d n_i + \lceil \log_2 b \rceil n_r \sum_{i=1}^d i \times n_i,$$

where n_r is the total number of rounds of flooding in the IGE process, which needs to be no less than $\max_i \{n_1, \dots, n_{d-1}\}$ to form a full-rank transmit power matrix. The last-hop nodes (at the d -th hop) do not affect n_r for having no downstream nodes. Since our approach conducts IGE together with flooding, the time consumption of IGE in our approach is $T_{our} = 0$.

For comparison, we also estimate the communication overhead for the traditional method of IGE, which asks nodes to take turns sending and listening measurement packets. In a multi-hop network, each node only needs to measure signals from upstream nodes. We can thus estimate the total communication overhead of the traditional method as

$$O_{other} = \lceil \log_2 b \rceil \sum_{i=1}^d n_i \times \sum_{j=0}^{i-1} n_j,$$

where $n_0 = 1$ presents for the initiator. The time consumption of the traditional method is equal to the number of total nodes (i.e., $T_{other} = \sum_{i=0}^d n_i$ slots). In a square network where each hop has an equal number of nodes (not including the initiator), we have $n_i = n_r = \sqrt{n}$, where n is the total number of nodes. The communication overhead of our and the traditional methods are $O_{our} = \lceil \log n_s \rceil nd + \lceil \log_2 b \rceil n \frac{d(d+1)}{2}$ and $O_{other} = \lceil \log_2 b \rceil n \frac{d(d+1)}{2}$, respectively. The extra communication overhead of our method is to construct the communication graph, which requires only $\lceil \log n_s \rceil$ bits for each node to record the reception time of the flooded packet. In contrast, our method needs no extra time for measurement, whereas the traditional method requires n extra time slots.

²The Nordic nRF52 series SoCs measure RSS with a granularity of 1dB.




Fig. 15. Testbed.

5 Performance Evaluation

This section presents and evaluates the implementation of our approach based on BlueFlood [30], a recent lower-power flooding protocol implemented on the Nordic nRF SoCs.

5.1 Experimental Setup

Implementation. We implement a multi-hop BLE network managed by an initiator. To support the centralized management, we modify BlueFlood to make each node capable of adjusting its configurations according to a scheduling plan generated by the initiator. This scheduling plan specifies the transmit powers of nodes for flooding and the power adjustments of nodes for IGE. In the IGE process, each node adjusts its transmit power in the specified time slot and uses the transmit power for flooding for the rest of the round. The number of rounds for IGE is set to the maximum number of nodes at all hops. On each node, we utilize the bit counter to take RSSI samples periodically for every byte received. The received power is estimated as the average of the measured RSSI samples. After the IGE process, nodes send back the received powers at different time slots to the initiator. We connect the initiator to a laptop for the IGE and power allocation.

Testbed and Configuration. We run all experiments in a 19-node testbed in our office, where the network topology is shown in Figure 15(b). The deployed network is of three hops with node 0 being the initiator. Each node can communicate with its neighbor-hop nodes and interfere cross-hop nodes. In our testbed, nodes at the first, second, and third hops are spatially aligned, making nodes prone to have similar channel gains to their neighbors and thus susceptible to destructive interference. Figure 15(c) shows the distribution of the channel gain similarities of nodes, measured by the differences between the two strongest channel gains for all the nodes, where 50% of nodes have the two strongest channels differing in less than 3dB. Our experiments are conducted on the Bluetooth advertising channel 37. For each node, when receiving a flooded packet, it rebroadcasts the packet for N_{Tx} times. As in BlueFlood, we set the total number of time slots in each round to $N_{Tx} + 2 \times \#TotalHops$, where $N_{Tx} = 4$ and $\#TotalHops$ is the total number of hops. In our experiments, we use the transmit powers in the linear region from $[-20, -16, -8, -4, 0]$ dBm.

5.2 Power Adjustment Selection

We want to select a proper power adjustment for IGE and understand its impact on network performance. As indicated by Equation (9), the error bound of IGE can be reduced by enforcing an identical, large power adjustment among all the nodes. To verify this, we repeat the experiments of flooding for more than 3,000 rounds under different power adjustments, where each IGE process takes 6 rounds followed by a 50-round flooding-alone period. Due to the limited available transmit powers, we cannot enforce all nodes to adopt the same power adjustment, but require that the power adjustments of nodes be no less than a threshold, denoted as ΔP . Before the start of experiments under each power adjustment, point-to-point channel estimation is conducted between




Fig. 16. The impact of power adjustment selection.

each pair of nodes to obtain the ground-truth channel gains. The channel gain (h) estimation error is computed by comparing the estimated channel gain from the power-domain IGE against the ground-truth ones. Figure 16 shows the impact of power adjustment selection on IGE and network performance. We can see from Figure 16(a) that the estimation errors of channel gains decrease significantly as the power adjustment ΔP increases. As a result, when the estimated channel gains are used to optimize the performance of flooding, the per-slot PER for flooding in the flooding-alone period also decreases, where the per-slot PER is the ratio of received packets over the number of valid receive slots. Figure 16(b) shows the impact of power adjustment on the IGE process. The stable per-slot PER is a result of two counter-balanced factors: increasing ΔP results in better channel estimation and power allocation but also causes higher interference. Although the per-slot PER in the IGE process remains stable under different Δ 's, the large power adjustment changes the packet loss distribution among nodes, with neighbors to the power-adjusted node having higher probability of receiving the flooded packets. This implies that when nodes at one hop adjust their transmit powers in sequence, using large power adjustment increases the likelihood of non-zero diagonal entries of the transmit power matrix at the next hop. We can see that the success probability of IGE increases sharply with ΔP from 53% to 95%. When an IGE process fails, the interference graph from the previous IGE process will be used. The probability of consecutive IGE failures is very low. In our experiments, we found no two consecutive failures of IGE. We use $\Delta P = 0.4mW$ for the following experiments.

5.3 Accuracy of Power-Domain IGE

To understand the accuracy of IGE, we compare the channel gains estimated from the power-domain IGE against the ground truth from the point-to-point channel estimation. The impact of channel variations is mitigated by conducting the power-domain and point-to-point IGEs back to back. We repeat this back-to-back experiment every 10 seconds for 60 times and collect more than 7,000 channel gains. Figure 17(a) shows the overall distribution of the estimation errors, where the estimation error is computed as the absolute difference between the estimated and the ground-truth channel gains. We can see that 68% of channels have an estimation error less than 4dB. Compared with the controlled experiments in Section 3.8, we find the estimation errors in real-world settings are in general larger due to varying and noisy network conditions. By further grouping channel gains by magnitude (see Figure 17(b)), we find that the estimation error increases as the channel gain decreases. The power-domain IGE excels at estimating large channel gains with the 75th percentile error less than 1.8dB for the channel gains greater than -40dB. For weak channels, since they have gains significantly smaller than the dominant channels, they have minimal influence on the resulting received power and thus are hard to be accurately estimated. It is worth noting that the estimation error of a channel in the power-domain IGE depends on whether adding this channel causes significant impact on the received power. Under the same transmit power, the




Fig. 17. Accuracy of IGE.

determining factor is the relative magnitude between strong and weak channels rather than the absolute magnitude. By considering for each node the channel of the strongest individually received power as the dominant channel and all the rest as the interference channels, we find that dominant channels all lie in $[-30, -20]$ dB, which can be accurately estimated with the 75th and 90th percentile errors less than 1.8dB and 1dB, respectively. This indicates that the power control algorithm correctly chose the strong channels as the dominant channels.

5.4 Performance Improvement of Flooding

To understand how IGE helps improve the performance of flooding, we want to compare our approach with BlueFlood. As our approach employs variable transmit powers and BlueFlood uses fixed transmit powers, it is difficult to make a fair comparison between them. Instead, we allow BlueFlood to use a higher average transmit power, equating to have a better link quality. Besides, we also develop a variant of BlueFlood with random transmit powers for comparison. In our experiments, we compare the performance of four schemes: (1) Optimized Flooding, optimizing the transmit powers of nodes in flooding with the estimated interference graph in the flooding-alone period; (2) Flooding+IGE, simultaneous flooding and IGE in the IGE process; (3) random-power BlueFlood, a variant of BlueFlood with random transmit powers; and (4) fixed-power BlueFlood, the original lower-power flooding protocol over BLE. We run the experiments of flooding for more than 3,000 rounds for each of the four schemes. Figure 18(a) shows the per-slot PERs under the four schemes, where the packet size is of 255 bytes. Due to the high channel similarity in our testbed (see Figure 15(c)) and the large packet size, all four schemes experience high PERs. In random-power BlueFlood, the channel similarity issue is mitigated using random transmit power such that it is easier for each node to have a channel of dominant received power. This increases the packet delivery rate of random-power BlueFlood, outperforming the fixed-power BlueFlood. The remaining two schemes, Optimized Flooding and Flooding+IGE, leverage the estimated interference graph to optimize the transmit powers of nodes for flooding and thus achieve lower per-slot PERs. Flooding+IGE performs worse than Optimized Flooding because it sacrifices performance for the purpose of IGE. As a result of the lower per-slot PER, Figure 18(b) shows that Optimized Flooding can achieve an **end-to-end (E2E)** PER of 2.4%, where the E2E PER is the ratio between the count of failed flooding rounds and the total round count. In contrast, the high per-slot PERs of the random- and fixed-power BlueFloods cause them to have E2E PERs as high as 8.1% and 13.4%, respectively.

Figure 18(c) shows the average latency for the four schemes in the unit of slots, where the latency is the time taken for all the nodes to receive the flooded packet and the error bars indicate the 10th and 90th percentile latencies, respectively. We can see that Optimized Flooding achieves the least latency among all the schemes and that it takes two fewer slots than the fixed-power BlueFlood on average to complete flooding. Meanwhile, even the 90th percentile latency of Optimized Flooding is no larger than the average latency of the fixed-power BlueFlood. Moreover, simultaneous flooding




Fig. 18. Performance improvement of flooding.

and IGE only slightly increases the average latency compared with Optimized Flooding. Looking into how nodes are covered over time during flooding (see Figure 18(d)), we find that, on average, Optimized Flooding can cover almost 90% of nodes in the first three slots, greatly improving the likelihood and speed of full coverage. Due to the optimized transmit powers of nodes in Optimized Flooding, nodes at each hop have a lower per-slot PER than other schemes. This explains why Optimized Flooding is more reliable and faster.

5.5 Power Consumption Overhead

Since IGE requires RSSI sampling and transmit power adjustment, we want to understand the power consumption overhead of these operations. We first set the experiment node in receive mode as in the Flooding+IGE scheme, which will use bit counter to collect RSSI samples for every byte received. We measure the working voltage and current of the node for at least 1,000 rounds with Keysight E36311A Power Supply on three devices: nRF52832 dongle, nRF52832 development kit (dev kit), and nRF52840 dev kit. For fair comparison, we repeat the experiment with BlueFlood (fixed power) on each device. Table 2 demonstrates the normalized average power consumption of BlueFlood and our IGE. Compared with BlueFlood, our IGE incurs only 1.60% power overhead on the dongle and no more than 3% on both development kits. To understand the overhead of transmit power adjustment, we enforced the nRF52 series devices to switch between different transmit powers for 10 million times and found that the process only took 3.75 seconds. This equates to 375ns per transmit power adjustment, which is three orders less than the slot time. We can thus safely neglect the power and time overhead caused by the transmit power switching.

Table 2. Normalized Power Consumption on Average

Board	nRF52832 Dongle	nRF52832 Dev Kit	nRF52840 Dev Kit
BlueFlood	100%	100%	100%
Our IGE	101.60%	102.28%	102.73%

6 Applications

To demonstrate the broader impact of IGE, we present two use cases for (1) convergecast and (2) channel map convergence in BLE networks. We consider a multi-hop BLE network with diverse network activities, including divergecast, convergecast, and distributed **peer-to-peer (P2P)** communication. The divergecast is implemented through concurrent flooding such that power-domain IGE can be conducted. It was shown previously that IGE helps improve the performance of flooding-based divergecast. In this section, we focus on using the interference graphs estimated from IGE to improve convergecast and distributed P2P communication in BLE networks.

6.1 Convergecast

Traditional Convergecast. Convergecast is a common network activity for data collection. We consider a typical tree topology for convergecast, where each parent node manages one or more child nodes and iterates through its child nodes to collect data. The data is collected hop by hop from the farthest leaf nodes to the root node, and each node transmits or receives data at scheduled time slots. Except for the root node, each non-leaf node acts as both a master and a slave. Parent and child nodes use **adaptive frequency hopping (AFH)** to mitigate external interference. Traditional AFH algorithms randomly select available channels from the channel map for frequency hopping [34]. Packet loss may occur when the channels of CTs collide.

IGE-Based Convergecast. IGE enables us to estimate the additive interference from multiple CTs on the same channel. Unlike traditional convergecast, IGE-based convergecast can identify links that can transmit concurrently on the same channel and those that cannot. We can avoid packet loss due to co-channel interference by allocating non-overlapping channels to mutually interfering links. To ease channel allocation, we put non-mutually interfering links in the same group and assign them the same set of channels. More specifically, we group links in a greedy manner: we first find the largest group of non-mutually interfering links from the ungrouped links and repeat the process until every link is in a group.

Experimental Settings and Results. We construct a tree topology connecting the root node and the rest of nodes based on the full network topology in Figure 15(b). Except for the leaf nodes, nodes with even indices are selected to be parent nodes. Specifically, node 0 manages all nodes in Hop 1, and each node with an even index in Hops 1 and 2 has two child nodes (i.e., node i is the parent of nodes $i + 5$ and $i + 6$). We consider the traditional convergecast as the baseline, where parent nodes randomly pick channels from the same set of available channels for channel hopping. For both IGE-based convergecast and baseline, no packet is retransmitted when packet loss occurs. Figure 19(a) shows the success probabilities of the two schemes for 10,000 repeated experiments, where each node has 37 available channels (namely all the BLE data transmission channels). It can be seen that IGE-based convergecast can effectively mitigate the impact of co-channel interference and achieve a success probability of 96.3%. This demonstrates that IGE succeeds in identifying mutually interfering links and allocates them non-overlapping channels. In contrast, the baseline only has a success probability of 82.6% due to co-channel interference. For failed convergecasts, we look into the percentage of nodes whose data are successfully collected. As shown in Figure 19(b), IGE-based convergecast can collect data from more nodes in the network than the baseline. Leveraging IGE can help convergecast avoid co-channel interference and reduce packet loss.




Fig. 19. Performance improvement of convergecast.

6.2 Channel Map Convergence

Traditional vs IGE-Assisted P2P Communication. In BLE networks, distributed P2P communications leverage frequency hopping to mitigate the impact of interference. For each transmitter-receiver pair, the available channels for frequency hopping are selected based on the channel map, which is dynamically updated by continuously monitoring the interference. A recent work [37] proposes to use the packet loss rate on each channel to update the channel map, where a channel is discarded if the packet loss rate on that channel exceeds a threshold. When multiple interferers exist, it generally takes long time for the channel maps of nodes to converge and packet loss occurs during the convergence process. With IGE, we can allocate non-overlapping channels to interfering links to both accelerate the convergence and reduce packet loss during the process. More specifically, we put links into groups such that links in the same group can transmit concurrently. We group links in a greedy manner based on the estimated interference graph, where the SINRs of nodes in each cluster exceed a threshold. We greedily find the maximum group for ungrouped links and repeat this for the rest of the ungrouped links until each link belongs to a group. For fairness, we evenly divide total available channels between groups.

Experimental Settings and Results. In our experiments, the initial channel map for each node includes 12 available channels. Traditional P2P directly employs AFH using the initial channel map, whereas the IGE-assisted P2P adjusts the channel maps of nodes based on the interference graph and employs AFH for data transmission using the adjusted channel maps. Figure 20(a) compares the **packet delivery rate (PDR)** between traditional and IGE-based P2Ps during the channel map convergence, where six transmitter-receiver pairs are formed with nodes at the second and third hops in our testbed. It can be seen that IGE-based P2P is more reliable than the traditional P2P and achieves a PDR of 99.9% even for the worst pair. Although traditional P2P also achieves a high average PDR, it suffers from severe sudden packet drops. As shown in Figure 20(b), there are three sudden jumps of packet loss, during which the packet loss rates are 8.7%, 11.3%, and 8.8%, respectively. This is because pairs 1 and 4 detect high packet loss rates on most available channels and have to reset their channel maps for a new round of convergence.

7 Discussions and Limitations

Non-Linearity Issues. We have identified four non-linearity issues found in our experiments and proposed several methods to control almost all the power ratios to fall between 0.9 and 1.1. However, there are still some outliers beyond this range. We ensure that these outliers are not measurement artifacts by excluding all possible known causes. Nonetheless, the limited visibility into the hardware of nRF52 series SoCs prevents us from diving deeper into this issue and further pushing the power ratios toward 1. The non-linearity issues weaken the linear relation between the transmit and received powers and affect the accuracy of IGE. Moreover, we limit the transmit powers of nodes within the linear region below 0dBm. Using a higher transmit power would greatly improve




Fig. 20. Performance improvement of channel map convergence.

the communication range. Based on the path-loss exponent model, when the path-loss exponent is 2, using the maximum transmit power (8dBm) of nRF52 series SoCs may improve the communication distance by 2.5 times than using a transmit power of 0dBm.

Scalability Issues. IGE requires constructing a full-rank transmit power matrix for nodes in the same hop. When the number of same-hop nodes is large, it takes a long time to construct the transmit power matrix. A longer IGE process does not delay flooding, as flooding is conducted simultaneously with IGE, but it slows down the update frequency of the interference graph, making it unable to follow the channel variations. For fast varying channels, the use of IGE is limited to small- to medium-sized networks.

Centralized Control, Overhead of IGE, and Optimizations. This work assumes a network with the needs of divergecast and convergecast, where a central node manages the network by collecting measurement data and disseminating decision plans. The integration of IGE introduces extra but small communication overhead to the network as analyzed in Section 4.4. We can reduce the overhead of IGE by (1) reducing the frequency of IGE or only conducting IGE when needed, (2) designing an efficient representation for control and measurement data, and (3) piggybacking these data onto normal traffic. Besides, estimating interference graphs also consumes computing resources. We have done this with a laptop connected to the central node, which is not necessary if efficient algorithms can be designed for the low-power devices. We may also consider decentralized power control for future optimizations, which alleviates both the communication and computing workload for the central node and is robust to the single point of failure.

Broader Impact of the Interference Graph. We have demonstrated how to achieve IGE together with CT-based flooding and improve the performance of flooding with the estimated interference graph. In fact, as interference graphs are widely used in wireless network resource management, the interference graph estimated from flooding can also be used to benefit many other network activities. Considering the multi-faceted contributions of IGE to the network, the related overhead is worthwhile.

8 Conclusion

In this article, we proposed to integrate IGE into data transmission tasks such that it can be done simultaneously with the data transmission tasks using the same frequency-time resources. We found that concurrent flooding is a perfect carrier for this integration in wireless sensor networks and designed network protocols and power allocation algorithms for efficient IGE. We showed that IGE on the Nordic nRF52 series SoCs was feasible in both controlled and real-world experiments and that flooding with power control based on the estimated interference graph outperformed the plain flooding. In addition to enhancing flooding, IGE was demonstrated to have potential for a broader range of tasks in BLE networks.

References

- [1] Beshr Al Nahas, Simon Duquennoy, and Olaf Landsiedel. 2019. Concurrent transmissions for multi-hop Bluetooth 5. In *Proceedings of the 2019 International Conference on Embedded Wireless Systems and Networks (EWSN '19)*. 130–141.
- [2] Kubra Alemdar, Divashree Varshney, Subhramoy Mohanti, Ufuk Muncuk, and Kaushik Chowdhury. 2021. RFClock: Timing, phase and frequency synchronization for distributed wireless networks. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (MobiCom '21)*. ACM, New York, NY, USA, 15–27. <https://doi.org/10.1145/3447993.3448623>
- [3] N. Baccour, A. Koubâa, L. Mottola, M. A. Zúñiga, H. Youssef, C. A. Boano, and M. Alves. 2012. Radio link quality estimation in wireless sensor networks: A survey. *ACM Transactions on Sensor Networks* 8, 4 (2012), 1–33.
- [4] M. Baddeley, C. A. Boano, A. Escobar-Molero, Y. Liu, X. Ma, U. Raza, K. Römer, M. Schuß, and A. Stanoev. 2020. The impact of the physical layer on the performance of concurrent transmissions. In *Proceedings of the IEEE 28th International Conference on Network Protocols (ICNP '20)*. IEEE, 1–12.
- [5] J. Cao, T. Peng, X. Liu, W. Dong, R. Duan, Y. Yuan, W. Wang, and S. Cui. 2019. Resource allocation for ultradense networks with machine-learning-based interference graph construction. *IEEE Internet of Things Journal* 7, 3 (2019), 2137–2151.
- [6] S. Chen, M. Coolbeth, H. Dinh, Y.-A. Kim, and B. Wang. 2009. Data collection with multiple sinks in wireless sensor networks. In *Wireless Algorithms, Systems, and Applications (WASA 2009)*. Lecture Notes in Computer Science, Vol. 5682. Springer, 284–294.
- [7] S. Chen, M. Huang, S. Tang, and Y. Wang. 2011. Capacity of data collection in arbitrary wireless sensor networks. *IEEE Transactions on Parallel and Distributed Systems* 23, 1 (2011), 52–60.
- [8] X. Chen and J. Huang. 2013. Distributed spectrum access with spatial reuse. *IEEE Journal on Selected Areas in Communications* 31, 1 (2013), 593–603.
- [9] W. Cui, K. Shen, and W. Yu. 2019. Spatial deep learning for wireless scheduling. *IEEE Journal on Selected Areas in Communications* 37, 6 (2019), 1248–1261.
- [10] H. Elgendi, M. Mäenpää, T. Levanen, T. Ihalainen, S. Nielsen, and M. Valkama. 2019. Interference measurement methods in 5G NR: Principles and performance. In *Proceedings of the 2019 16th International Symposium on Wireless Communication Systems (ISWCS '19)*. IEEE, 233–238. <https://doi.org/10.1109/ISWCS.2019.8877215>
- [11] S. C. Ergen and P. Varaiya. 2010. TDMA scheduling algorithms for wireless sensor networks. *Wireless Networks* 16, 4 (2010), 985–997. <https://doi.org/10.1007/s11276-009-0183-0>
- [12] F. Fernandes, A. Ashikhmin, and T. L. Marzetta. 2013. Inter-cell interference in noncooperative TDD large scale antenna systems. *IEEE Journal on Selected Areas in Communications* 31, 2 (2013), 192–201.
- [13] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh. 2011. Efficient network flooding and time synchronization with Glossy. In *Proceedings of the Conference on Information Processing in Sensor Networks (IPSN '11)*. IEEE, 73–84.
- [14] M. M. Halldörsson. 2012. Wireless scheduling with power control. *ACM Transactions on Algorithms* 9, 1 (2012), 1–20.
- [15] J. Harbin, A. Burns, R. I. Davis, L. S. Indrusiak, I. Bate, and D. Griffin. 2019. The airtight protocol for mixed criticality wireless CPS. *ACM Transactions on Cyber-Physical Systems* 4, 2 (2019), 1–28.
- [16] Carsten Herrmann and Marco Zimmerling. 2023. RSSISPY: Inspecting concurrent transmissions in the wild. In *Proceedings of the 2022 International Conference on Embedded Wireless Systems and Networks (EWSN '22)*. ACM, New York, NY, USA, 48–59.
- [17] Aravind Iyer, Catherine Rosenberg, and Aditya Karnik. 2006. What is the right model for wireless channel interference? In *Proceedings of the 3rd International Conference on Quality of Service in Heterogeneous Wired/Wireless Networks (QShine '06)*. ACM, New York, NY, USA, 2-es. <https://doi.org/10.1145/1185373.1185376>
- [18] R. Jacob, A. B. Schaper, A. Biri, R. Da Forno, and L. Thiele. 2020. Synchronous transmissions on Bluetooth 5 and IEEE 802.15.4—A replication study. In *Proceedings of the 3rd Workshop on Benchmarking Cyber-Physical Systems and Internet of Things (CPS-IoTBench '20)*. ACM, New York, NY, USA, 1–39. <https://doi.org/10.1145/3604430>
- [19] Haifeng Jia, Yichen Wei, Zhan Wang, Jiani Jin, and Yibo Pi. 2023. Efficient interference graph estimation via concurrent flooding. In *Proceedings of the 2023 International Conference on Embedded Wireless Systems and Networks (EWSN '23)*. ACM, New York, NY, USA, 249–255.
- [20] S. Kurt and B. Tavli. 2017. Path-loss modeling for wireless sensor networks: A review of models and comparative evaluations. *IEEE Antennas and Propagation Magazine* 59, 1 (2017), 18–37.
- [21] M. Lee, G. Yu, and G. Y. Li. 2020. Graph embedding-based wireless link scheduling with few training samples. *IEEE Transactions on Wireless Communications* 20, 4 (2020), 2282–2294.
- [22] X.-Y. Li, K. Moaveni-Nejad, W.-Z. Song, and W.-Z. Wang. 2005. Interference-aware topology control for wireless sensor networks. In *Proceedings of the 2005 2nd Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks (SECON '05)*. IEEE, 263–274. <https://doi.org/10.1109/SAHCN.2005.1557081>

- [23] C. H. Liao, Y. Katsumata, M. Suzuki, and H. Morikawa. 2016. Revisiting the so-called constructive interference in concurrent transmission. In *Proceedings of the 2016 IEEE 41st Conference on Local Computer Networks (LCN '16)*. IEEE, 280–288.
- [24] Z. Q. Luo and P. Tseng. 1994. Perturbation analysis of a condition number for linear systems. *SIAM Journal on Matrix Analysis and Applications* 15, 2 (1994), 636–660.
- [25] Ritesh Maheshwari, Shweta Jain, and Samir R. Das. 2008. A measurement study of interference modeling and scheduling in low-power wireless networks. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems (SenSys '08)*. ACM, New York, NY, USA, 141–154. <https://doi.org/10.1145/1460412.1460427>
- [26] Ritesh Maheshwari, Shweta Jain, and Samir R. Das. 2008. On estimating joint interference for concurrent packet transmissions in low power wireless networks. In *Proceedings of the 3rd ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization (WiNTECH '08)*. ACM, New York, NY, USA, 89–94. <https://doi.org/10.1145/1410077.1410094>
- [27] R. Maheshwari, S. Jain, and S. R. Das. 2009. A measurement study of interference modeling and scheduling in low-power wireless networks. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems (SenSys '09)*. ACM, New York, NY, USA, 141–154. <https://doi.org/10.1145/1460412.1460427>
- [28] E. Main and D. Coffing. 2022. An FSK demodulator for Bluetooth applications having no external components. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing* 49, 6 (2022), 373–378.
- [29] A. U. Makarfi, K. M. Rabie, O. Kaiwartya, K. Adhikari, G. Naurzybayev, X. Li, and R. Kharel. 2020. Toward physical-layer security for Internet of Vehicles: Interference-aware modeling. *IEEE Internet of Things Journal* 8, 1 (2020), 443–457.
- [30] Beshr Al Nahas, Antonio Escobar-Molero, Jirka Klaue, Simon Duquennoy, and Olaf Landsiedel. 2021. BlueFlood: Concurrent transmissions for multi-hop Bluetooth 5—Modeling and evaluation. *ACM Transactions on Internet of Things* 2, 4 (2021), Article 22, 30 pages. <https://doi.org/10.1145/3462755>
- [31] Nordic Semiconductor. 2021. *nRF52832 Product Specification v1.8*. Nordic Semiconductor, Trondheim, Norway.
- [32] Nordic Semiconductor. 2021. *nRF52840 Product Specification v1.7*. Nordic Semiconductor, Trondheim, Norway.
- [33] E. Peralta, G. Pocovi, L. Kuru, K. Jayasinghe, and M. Valkama. 2022. Outer loop link adaptation enhancements for ultra reliable low latency communications in 5G. In *Proceedings of the 2022 IEEE 95th Vehicular Technology Conference (VTC2022-Spring '22)*. IEEE, 1–7.
- [34] Bluetooth SIG. 2023. *Bluetooth Core Specification*. Bluetooth SIG. <https://www.bluetooth.com/specifications/specs/core-specification-5-4/>
- [35] Dongjin Son, Bhaskar Krishnamachari, and John Heidemann. 2006. Experimental study of concurrent transmission in wireless sensor networks. In *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems (SenSys '06)*. ACM, New York, NY, USA, 237–250. <https://doi.org/10.1145/1182807.1182831>
- [36] D. Son, B. Krishnamachari, and J. Heidemann. 2006. Experimental study of concurrent transmission in wireless sensor networks. In *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems (SenSys '06)*. ACM, New York, NY, USA, 237–250. <https://doi.org/10.1145/1182807.1182831>
- [37] Michael Spörk, Jiska Classen, Carlo Alberto Boano, Matthias Hollick, and Kay Römer. 2020. Improving the reliability of Bluetooth Low Energy connections. In *Proceedings of the 2020 International Conference on Embedded Wireless Systems and Networks (EWSN '20)*. 144–155.
- [38] Q. Tang, L. Yang, G. B. Giannakis, and T. Qin. 2007. Battery power efficiency of PPM and FSK in wireless sensor networks. *IEEE Transactions on Wireless Communications* 6, 4 (2007), 1308–1319.
- [39] Y. Yu, E. Dutkiewicz, X. Huang, and M. Mueck. 2013. Downlink resource allocation for next generation wireless networks with inter-cell interference. *IEEE Transactions on Wireless Communications* 12, 4 (2013), 1783–1793.
- [40] G. Zhou, T. He, J. A. Stankovic, and T. Abdelzaher. 2005. RID: Radio interference detection in wireless sensor networks. In *Proceedings of the IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies*. IEEE, 891–901.

Received 5 December 2023; revised 7 June 2024; accepted 11 August 2024