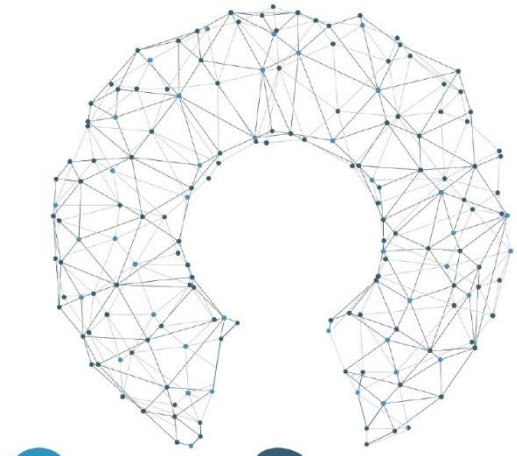

www.openswarm.eu

Call: HORIZON-CL4-2022-DATA-01

Type of action: RIA

Grant agreement: 101093046

Deliverable n° : D2.2 Zero-Wire Concept Translated from Optical to RF-based Communication

Work Package n°: WP2: Orchestration of Collaborative Smart Nodes

Task Lead: ADI

WP Lead: ADI

This project has received funding from the European Union's Horizon Europe Framework Programme under Grant Agreement No. 101093046.

Funded by the European Union

Document information	
Author(s)	Danny Hughes, Jeroen Famaey, Xinlei Liu, Andrey Belogaev, Jonathan Oostvogels, Mengyao Liu, Bingwu Fang
Reviewers	Ritesh Singh, Thomas Watteyne
Submission date	31-Oct-2024
Due date	31-Oct-2024
Type	Report
Dissemination level	PU

Document history			
Date	Version	Author(s)	Comments
11-Jun-2024	01	Danny Hughes	preview
31-Oct-2024	02	Danny Hughes	deliverable

Executive Summary

This deliverable summarizes the technical outcomes of Task 2.2 of Work Package WP2. The goal of the task is to port the symbol synchronous communication paradigm of ZeroWire [1] from optical to Radio Frequency (RF). This resulted in three papers, each exploring this approach at a different position in the frequency spectrum:

- **mmWave ZeroWire Analog** (Asynchronous Burst Link - ABL) was realized in simulation using fast pulse modulation at 60 GHz. ABL delivers 99% reliability and under 1 ms latency for 32b packets relayed over 8 hops and a data rate of 64 kbps. This work is published in [2].
- **ISM ZeroWire Analog** was realized in simulation, also uses fast pulse modulation at 2.4 GHz, achieving latency of a few milliseconds over multiple hops with high reliability, with only a 0.3% latency increase per extra hop and a data rate of 100 kbps. It is published in [3].
- **VHF ZeroWire Analog** (Capacitive and Inductive Network - CaIN) was realized as a hardware/software reference design. CaIN uses Near Field Communication (NFC) at 40MHz to deliver a novel transceiver that lowers latency by eliminating Radio Duty Cycling (RDC). CaIN delivers 95% reliability, sub-µW receiver power, 20kbps goodput and under 50ms latency for a 10-hop network [4].

Considered in sum, these three research initiatives thoroughly investigate how symbol synchronous communication can be brought to RF media. All experimental approaches significantly improve latency compared the state of the art and provide a starting point for further research across the frequency spectrum.

1. Introduction

Work Package 2 (WP2) aims to develop the core technologies that are required to orchestrate collaborative smart nodes, including networking, security, coordination and management. In this context, current Internet of Things (IoT) mesh networks, such as Bluetooth Low Energy, WirelessHART or 6TiSCH have emerged as a scalable solution for embedded networks in large and/or cluttered environments. By implementing routing functionality on top of short-range radios, these networks realize an extensible wireless fabric that can cover challenging physical environments. WMN technologies have a proven track record of delivering high end-to-end reliability while preserving multi- year battery life in power-constrained environments [5, 6, 7].

Current mesh networks provide reliable multi-hop communication but are poorly suited for swarm computing due to their **high and non-deterministic latency** [1], For example, Bluetooth Mesh latency may be as high as multiple seconds, with double digit rates of packet loss [8]. Wireless networks for Industrial Automation for Process Automation (WIA-PA) exhibit near-zero error rates, but still face latency in the tens of milliseconds [9]. WirelessHART, SmartMesh IP and 6TiSCH may be designed for real-time process monitoring and control, delivering high reliability with low error rates, but, these approaches likewise fail to address end-to-end latency challenges.: various application domains require millisecond or sub-millisecond end-to-end latency [10], This deliverable details Task 2.2, which tackles this lack of low-latency mesh networking support by developing modulation schemes that are better suited for integration in low-latency densely connected systems, at a variety of RF frequencies (VHF, ISM* and mmWave).

The root cause for the high latency faced by existing IoT mesh networks is that they primarily adopt modulation schemes that require store-and-forward networking strategies, thus necessitating nodes to fully receive a data frame before forwarding it across a mesh. Such strategies fall short in meeting the growing demands of applications that require both high reliability and low latency concurrently. Exploring

alternatives to traditional store-and-forward strategies and collision avoidance protocols is essential to effectively fulfill these demands, but the performance of such alternatives heavily depends on physical-layer properties [11].

Prior to the work described in this document, Zero-Wire [1], introduced by KU Leuven in 2020, addressed the low-latency mesh networking problem by delivering *symbol-synchronous* transmission, wherein distributed nodes collaborate in the transmission of the smallest data unit – the symbol – across the network illustrates symbol synchronous transmission in comparison to store-and-forward and concurrent approaches.

The original ZeroWire approach results in up to 19 kbps of contention-agnostic goodput, latency under 1 ms for two-byte packets, and jitter under 50 µs. However, the current Zero-Wire prototype is realized using optical communication and cannot be applied directly to RF due to the non-additive nature of RF signals, which may interfere constructively or destructively when transmitted in parallel. Moreover, Zero-Wire does not consider the energy constraints faced in wireless cyber-physical systems. The work done for Task 2.2 addresses these gaps.

The following chapters detail the research performed for Task 2.2. Chapter 2 further positions our contributions relative to the state of the art and provides background information necessary to understand the concepts relied on throughout this document. The research done for Task 2.2 has resulted in results published papers and a third paper under submission, which are presented in Chapters 3 through 5. Chapter 6 concludes this document.

2. Key Performance Indicators (KPIs)

This task concerns *Scientific Objective SO1: Orchestration of Collaborative Smart Nodes* and specifically *communication*. The KPIs for low-power wireless networking focus on end-to-end reliability, throughput, latency and power consumption.

Table 2.1: Key Performance Indicators. These KPIs summarize the results presented in this document, i.e. they complement the more high-level KPIs defined in the Grant Agreement DoA (shown in bold).

KPI	ABL	ISM ZeroWire	CaIN
Reliability	99%	99%	95%
Link Layer Throughput	64kbps	100kbps	20kbps
Multi-Hop Latency	< 1ms	< 1ms	< 50ms over 10 hops.
Power	NA (Simulation)	NA (Simulation)	20mW transmit < 1µW receive

As can be seen from **Table 2.1**, the technologies described in this deliverable offer a compelling range of KPIs, with demonstrated or simulated reliability between 95% and 99%, throughput in the 10s of kbps and low latency even in the case of large multi-hop networks. In terms of latency, all technologies demonstrate large markups in comparison to conventional mesh networks (> 100x). In terms of power consumption, CaIN achieves a ~5000x improvement in listening power consumption in comparison to state-of-the-art transceivers such as SmartMesh IP and Bluetooth Low Energy. This enables the CaIN transceiver to remain in an always-on and therefore low latency mode.

Reliability is lower than mature solutions such as SmartMesh IP, which guarantees 99.99% end-to-reliability. However, this stream of research drives latency and power consumption significantly lower.

3. Technology Readiness Level

Before the start of the project, the symbol-synchronous bus paradigm was at TRL 2, with radio-frequency instances of the paradigm at TRL 1. The technologies described in this document advance the TRL in several regions of spectrum: **Table 3.1** summarizes the TRL level of the three core technologies described in this deliverable. All technologies were developed during the course of the project.

Table 3.1: Technology Readiness Level

	TRL Level	Description
ABL	**TRL 2**	Analytical studies and initial validation have been performed through simulation.
ISM ZeroWire	**TRL 2**	Analytical studies and initial validation have been performed through simulation.
CaIN	**TRL 5**	An open-source pre-prototype has been developed (hardware + software) and tested in the lab at small scale (networks of up to 10 devices).

4. Background

In this section, we give a brief introduction of store-and-forward technology, followed by a discussion of non-store-and-forward technology, that is, synchronous (or "concurrent") and symbol synchronous transmission.

Store-and-Forward: In wireless mesh networks, reliability is challenged by collisions and interference. Techniques such as Carrier Sense Multiple Access (CSMA) and Time Synchronised Channel Hopping (TSCH) mitigate these issues by minimizing collisions and re-transmissions, while store-and-forward routing preserves data integrity and facilitates orderly traffic flow. Together, these strategies synergize to provide reliable network service [11]. However, these approaches achieve reliable networking at the expense of latency, as illustrated in Figure 4.1.

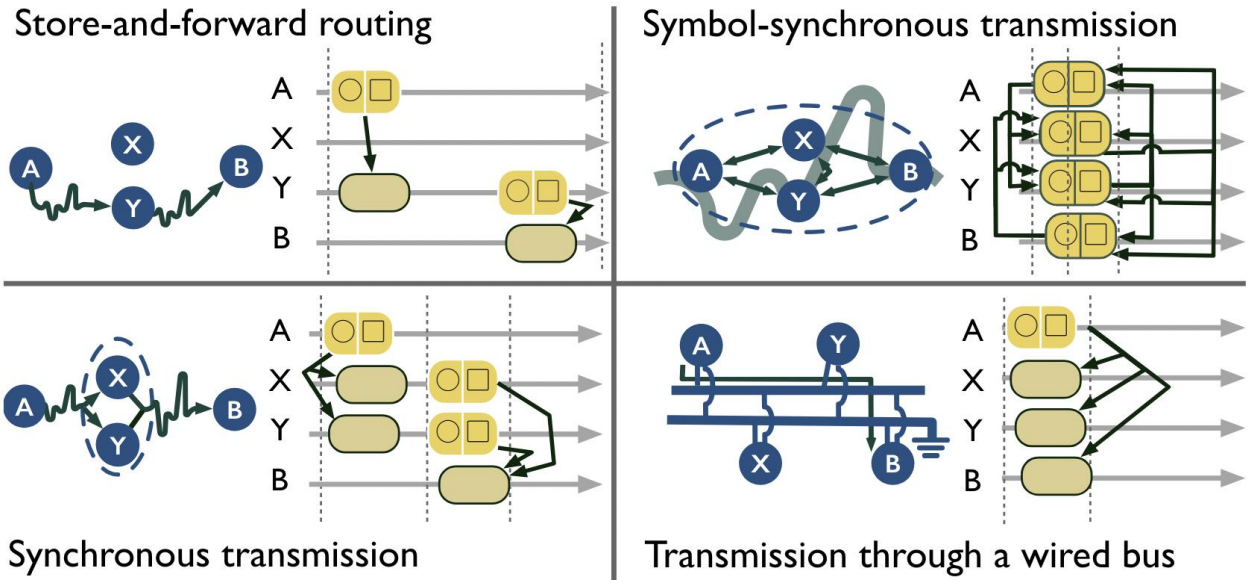

Figure 4.1: Symbol-synchronous transmission behaves like a wired bus avoiding the latency incurred by store-and-forward or synchronous transmission.

Protocols based on store-and-forward require nodes to sequentially receive, check, and forward data. Concretely, node A senses the channel state to prevent collisions and/or schedules data transmissions to node Y based on a routing strategy. Once the packet is fully received, node Y, in turn, assesses the channel state and/or waits for a transmit slot before forwarding the packet according to the routing policy and so on to the subsequent destination. While reliable store-and-forward networking is widely available, the necessity of fully receiving packets, monitoring channel states, checking routing information, and waiting for a transmission slot results in high E2E latency that scales quickly with the number of hops [1].

Concurrent Transmission: Concurrent Transmission protocols allow multiple nodes to broadcast packets simultaneously [12]. Thanks to the capture effect and non-destructive interference, nodes are likely to receive at least one of these simultaneous transmitted signals correctly [11].

Glossy, a pioneering concurrent transmission approach, achieves a few milliseconds latency and 99.99% reliability [12]. Glossy is a typical network flooding architecture targeting one-to-all communication across a WSN [13]. Glossy proposes to design a time-synchronization network to exploit the constructive interference for decoding the packet successfully. In Glossy, an initiator starts the broadcast process in a half-duplex way. Once the one-hop neighbors of the initiator have received signal, they re-broadcast in close synchrony. The signal is re-transmitted in a broadcast way until all nodes receive it successfully. Blueflood is another packet synchronous dissemination protocol performing similarly to Glossy, but it is based on a multi-hop Bluetooth mesh network [15]. In addition, a novel routing scheme called interference coordinated routing (ICR) is designed to make use of constructive interference and concurrent transmission to decrease the latency.

Recent works have demonstrated the versatility of concurrent transmission mechanisms, extending beyond their initial application in IEEE 802.15.4 [12]. This approach is now effectively implemented in various technologies, including Ultra-WideBand (UWB) [14], Bluetooth Low Energy (BLE) [15], and Long Range (LoRa) [16].

These advancements underscore the adaptability of concurrent transmissions across different physical layers.

Although all the above protocols benefit from concurrent transmissions, they still transmit the information packet by packet, which forces every relay node to wait before it receives and decodes all symbols included in a packet before it can relay this packet. This waiting time imposes a negative impact on the end-to-end latency of a multi-hop network since the latency linearly increases by the whole packet transmission time with each hop. For instance, in Glossy, for a 64-bit packet, the end to-end latency increases more than twice from 0.81ms to 1.77ms as the number of hops increases from 3 to 5 [13].

Symbol Synchronous Transmission: Prior research on concurrent transmissions has primarily focused on packet-wise operations, which naturally results in packet-scale latency [12, 15, 16], quantified on at least millisecond level. To overcome such poor scaling behavior, a technology supporting to forward data symbol-per-symbol instead of packet-per-packet is proposed, to avoid a linear increase of the latency as a function of a number of hops. Zippy is a transmission protocol for one-to-all communication that exploits symbol-per-symbol transmissions. Its advance is to design an asynchronous wake-up strategy for every node and exploit a simple On-off keying (OOK) transmitter to achieve per-hop synchronous transmission [17]. The Zippy network encodes every symbol using repetition codes and uses carrier frequency randomization to avoid destructive interference. It achieves deterministic performance with an end-to-end latency of tens of microseconds and low power consumption. However, Zippy only performs well in a network with a maximum of two hops, while its performance crucially degrades for the many-hop network and long distance communication. Its scalability is also limited due to the usage of different carrier frequencies. Zero-Wire [1], however, introduces a novel approach by implementing a symbol-synchronous bus network protocol that operates at the level of the smallest unit of transmission (i.e. the symbol), allowing for dramatically lower multi-hop latency.

This mechanism allows relay nodes to immediately forward symbols upon detection, without coordinating with others, and without making transmission decisions based on the contents of the signal, thus achieving sub-millisecond E2E latency. In symbol

synchronous transmission, nodes relay incoming symbols with a relay time offset D_{offset}, such that the time difference D_{jitter} between the moment the first instance of a symbol originally transmitted by the initiator is detected ($T_{firstArrival}$) and the detection of the last instance ($T_{latestArrival}$), due to different relay paths, is smaller than the symbol duration, as formalised below.

$$D_{jitter} = T_{latestArrival} - T_{firstArrival} \quad (1)$$

$$D_{jitter} < D_{symbol} \quad (2)$$

Under condition (2), relayed signals are not re-relayed again if their initial reception and subsequent relayed version are observed with a time offset less than the symbol duration. This ensures effective symbol propagation across one-hop neighbors, effectively preventing a cycle of re-relaying identical symbols by the same nodes.

Note that [1] was not done as part of the OpenSwarm project, but rather was the foundation for the addition of ZeroWire in the proposal.

5. mmWave ZeroWire Analog (30-60GHz)

This section describes a mmWave ZeroWire-like protocol realized in simulation at 60GHz. The text is based on [2].

The emergence of millimeter-wave (mm-wave) radio front-ends, targeting carrier frequencies of 30GHz to 300GHz presents a promising approach to bring ZeroWire-like communication to radio space due to the relative abundance of bandwidth, Achieving this vision in practice however, requires careful design of mm-wave waveforms and symbol-synchronous modulation logic that take into account the realities of signal

generation and propagation at those frequencies. To tackle this challenge, we propose Asynchronous Burst Link (ABL), an innovative multi-hop mesh network architecture which employs symbol-synchronous transmission of mm-wave pulses. Our evaluation of ABL involves a simulated 9-node network with a maximum of 8 hops, utilizing 32-bit packets. The results indicate that latency for the most distant node remains under 1 millisecond, while maintaining 99% reliability. Our findings present evidence for the efficacy of novel mm-wave transceiver architectures in meeting the low latency and robustness demands of network services crucial for critical industrial applications.

Additionally, mm-wave technology, operating within the 30-300GHz spectrum, emerges as a key solution to effectively tackle the growing demand for timely data transmission. With extensive spectral resources, mm-wave enables multi-gigabit-per-second data rates, a vital aspect in contemporary high-speed communication networks [18]. The application of mm-wave technology in the wireless control of insulated gate bipolar transistors (IGBTs) serves as a prime example of its potential, attaining approximately 400 ns ultra-low latency [19]. However, this is for a point-to-point communication model and, given the problematic propagation of mm-wave, this is unsuitable for covering large and complex areas, where mesh networking is typically required.

5.1 Introduction

In response to the necessity of low latency mesh networking, we propose a novel mm-wave mesh network paradigm; Asynchronous Burst Link (ABL). Our main contributions are as summarized as follows:

- We introduce the first *symbol-synchronous* mm-wave radio protocol, a technique heretofore exclusive to Optical Wireless Communication (OWC) [1]. By employing a symbol-synchronous approach, one-hop neighbors selectively and quickly relay signals corresponding to identical bits without feedback loops or channel estimation. This novel integration effectively reduces latency and simplifies meshed system architecture. Additionally, the multi-hop mesh architecture of ABL effectively extends the coverage of mm-wave communications.

- We introduce a novel modulation scheme, Differential Pulse (DP) modulation, which ensures reliable network service and addresses errors caused by destructive interference between concurrent signal transmissions, as well as mitigates environmental background noise.

These contributions are simulated using MATLAB R2023b with the Communications Toolbox, and their evaluation demonstrates promising results, including sub-millisecond latency with 99% reliability across a 9-node mesh network featuring a maximum of 8 hops, employing 32-bit packets. The lack of a time-synchronisation requirement between ABL nodes prior to data exchange is a perfect fit with recent initiatives in the area of crystal-free radios, such as the Single Chip Micro Mote (SCμM) [20]. Prototype SCμM motes are now available with mm-wave.

5.2 Design

This section describes the design of ABL starting from the proposed Differential Pulse (DP) Modulation scheme followed by the comprehensive ABL protocol. Finally, we describe how the system can be integrated with conventional mm-wave hardware.

Differential Pulse Modulation: In Zero-Wire [1], optical signals are propagated symbol-synchronously, with simultaneous transmission of identical symbols by one-hop neighbors. This system employs On-Off Keying (OOK) for data encoding, modulating ambient light intensity levels. Unlike the optical media of the original Zero-Wire systems, radio-based systems require techniques to address destructive interference which occurs due to the use of coherent transmitters in conventional radios. To mitigate these problems, it is essential to develop an efficient modulation and demodulation scheme that circumvents destructive interference through redundancy in time and frequency space.

Inspired by Time Spread On-Off Keying (TS-OOK) modulation with pulse signals in Terahertz communication [21], we propose the Differential Pulse (DP) modulation and demodulation scheme to transmit data in narrow pulses to mitigate the potential destructive interference. The details are as follows.

Modulation: In the DP modulation scheme, binary information is encoded through the variation in maximum amplitude as depicted in **Figure 5.1**.

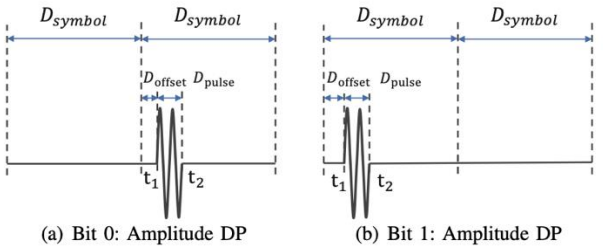

Figure 5.1: DP Modulation scheme for 0/1 binary information representation.

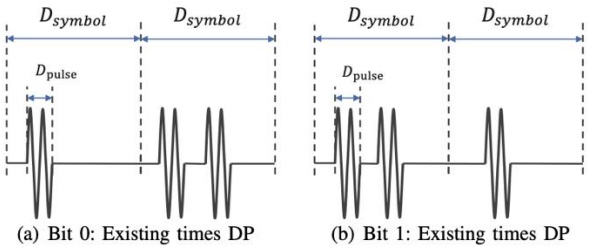

Figure 5.2: DP Demodulation scheme for 0/1 binary information representation.

The source of the signal, or initiator, encodes binary information by toggling a sine wave signal generator On and Off at predetermined intervals. Concretely, with the proposed DP scheme, bit-1 is encoded by a "pulse exists" symbol followed by a "silence" symbol, and vice versa for a bit-0, as demonstrated in **Figure. 5.1(a)** and **Figure 5.1(b)**. Note that in real systems, the symbol duration may be many times longer than the pulse duration, which has here been tailored for visual clarity. The corresponding representation of 0/1 binary information is as follows:

$$S^0 = \begin{cases} A \cdot \sin\left[2\pi \cdot f_c \cdot (t - D_{offset}) + \phi\right] \\ \text{when } 0 \le t_1 < t < t_2 \le D_{symbol}, \\ 0 \\ \text{when } \{t | t \in [0, 2 \cdot D_{symbol}] \, \Lambda t \notin [t_1, t_2]\} \end{cases}$$

$$S^1 = \begin{cases} A \cdot \sin\left[2\pi \cdot f_c \cdot t + \phi\right] \\ \text{when } D_{symbol} \le t_1 < t < t_2 \le 2 \cdot D_{symbol}, \\ 0 \\ \text{when } \{t | t \in [0, 2 \cdot D_{symbol}] \, \Lambda t \notin [t_1, t_2]\} \end{cases}$$

where A is the amplitude of the pulse and fc is the frequency of the sine wave pulse. The timing offset, denoted as D_{offset}, arises due to hardware imperfections, φ is the phase shift of the signal, $\varphi \in [0, 2\pi]$. D_{symbol} is the duration of symbol. The duration of a pulse is $D_{pulse} = t_2 - t_1$. It is much smaller than the symbol duration, i.e., $D_{pulse} \ll D_{symbol}$. The relationship between D_{pulse} and D_{symbol} can be simply illustrated as, $D_{symbol} = \beta \cdot D_{pulse}$ where $\beta > 0$. In our proposed DP scheme, 1-bit data is encoded using a pair of consecutive symbols, thereby, the duration for 1-bit binary information is $D_{bit} = 2 * D_{symbol}$.

Demodulation: As depicted in Figure. 5.1(a), a signal is decoded as bit-0 if the amplitude of the pulse signal in the first symbol is less than in the subsequent symbol, and vice versa for bit-1, as shown in Figure. 5.1(b). Moreover, as shown in Figure 5.2(a), a signal is decoded as bit-0 if the maximum time of pulse signal exist in the first symbol is shorter than that in the following symbol pulse, and vice versa for bit-1, as shown in Figure 5.2(b). If Fa > Ft, the demodulation scheme base on the amplitude difference, otherwise base on the maximum time difference. Concretely:

$$bit-0=\begin{cases}\max\left(A_{pulse}^{1st}\right)<\max\left(A_{pulse}^{2nd}\right)\\ \max\left(T_{pulse}^{1st}\right)<\max\left(T_{pulse}^{2nd}\right)\end{cases}$$

$$bit-1=\begin{cases}\max\left(A_{pulse}^{1st}\right)>\max\left(A_{pulse}^{2nd}\right)\\ \max\left(T_{pulse}^{1st}\right)>\max\left(T_{pulse}^{2nd}\right)\end{cases}$$

$$F_a=abs(A_{pulse}^{1st}-A_{pulse}^{2nd})/min(A_{pulse}^{1st},A_{pulse}^{2nd})$$

$$F_t=abs(T_{pulse}^{1st}-T_{pulse}^{2nd})/min(T_{pulse}^{1st},T_{pulse}^{2nd})$$

ABL Protocol: In our ABL system, nodes continuously sample and detect the amplitude of incoming signals to decide whether to re-transmit. If the amplitude is higher than a predefined threshold, a node re-transmits the incoming pulse with a slight time offset. If the amplitude of signal falls below the predefined threshold, and thus should not be relayed, nodes still attempt to decode the signal. The corresponding procedure is illustrated in Figure 5.3.

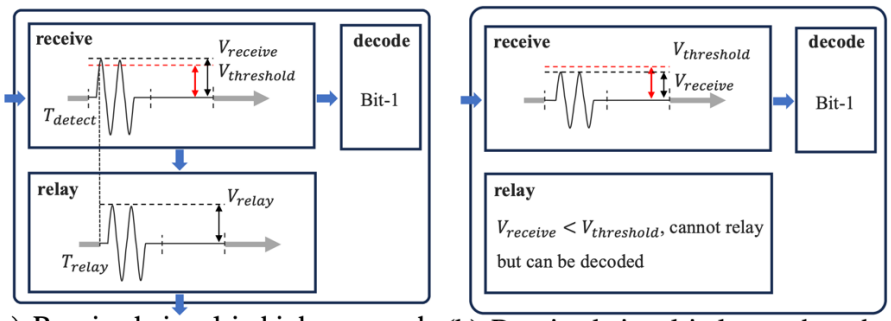

(a) Received signal is higher enough to be detected

(b) Received signal is lower than the threshold, but can be decoded

Figure 5.3: Signal receiving, detection, relaying, and decoding.

As shown in Figure 5.3, a single node includes three main functions, receiving, relaying base on the signal detection, and decoding according to our proposed signal decoding method. We denote $V_{receive}$ as the amplitude of the received signal, and $V_{threshold}$ as the system's signal detection threshold, configured based upon the noise floor. As shown in Figure 5.3(a), if the received signal is larger than $V_{threshold}$, a node starts to relay the signal at detection time T_{detect}. If the amplitude of the received signal is not large enough to be detected, but not been fully overwhelmed by the noise, nodes still attempt to decode the signal and get the transmitted data in time, as depicted in Figure 5.3(b).

ABL architecture A potential low-level (LL) fine-grained system architecture for ABL is shown in Figure 5.4(a). On the transmitter side, a Voltage-Controlled Oscillator (VCO), such as 840UF- 29/3830 [22], is employed to generate the 60GHz mm-wave signal. The RF-Switch is controlled by a microcontroller, facilitating the creation of pulse signals as well as the configuration of pulse duration. These signals are subsequently amplified and transmitted via an antenna. Conversely, on the receiver side, incoming signals captured by the antenna are initially amplified by a Low-Noise Amplifier (LNA). The amplified signals are then digitized by an Analog-to-Digital Converter (ADC), followed by signal processing to decide whether to retransmit.

The ABL system can also be implemented with existing commercial off-the-shelf (COTS) hardware products on top of a higher-level (HL) modulation method, such as with TI mm-wave radars [23]. Concretely, the corresponding potential hardware architecture is shown in Figure 5.4(b).

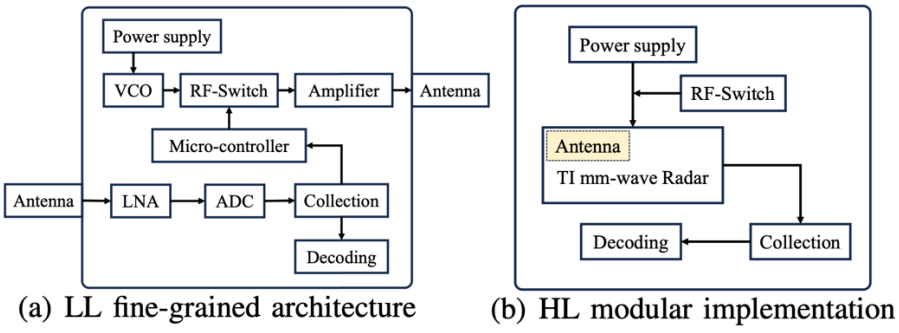

Figure 5.4: ABL System Architecture.

5.3 Evaluation

The proposed mechanisms are simulated in Matlab on a 8-core 2.3GHz machine with an 32-GB RAM. In this section, we experimentally analyse latency, reliability and robustness.

Table 5.1: Network Configuration

Parameter	Value	Note
Pt	20 dBm	Transmission power
fc	60 GHz	Carrier frequency
fs	25 MHz	Sampling frequency
Dpulse	200 ns	Pulse duration
Dsymbol	1µs	Symbol duration
Lpacket	32 bit	Packet length

Simulation Set-up: To ensure that our simulation stays close to achievable results, we parameterize our simulation to match current technological standards and hardware platforms. Simulation parameters are configured as shown in **Table 5.1** and explained in the following text.

In our simulation, we implement a Ray-Tracing propagation model, integrating Additive White Gaussian Noise (AWGN). Our focus is solely on Line-of-Sight communication, assuming an environment free of blockages. As shown in **Figure 5.5**, our setup includes 9 nodes, with E1(0,0) designated as the initiator of transmission. Nodes communicate only with their immediate one-hop neighbors due to signal strength limitations. For instance, as shown in **Figure 5.5**, the yellow lines between E1, and E2, E4 indicate that E1 communicates with E2 and E4, but not with E5, E3, or E7. When direct communication between the initiator and destination fails, symbols follow a multi-hop route of between 2 and 8 hops.

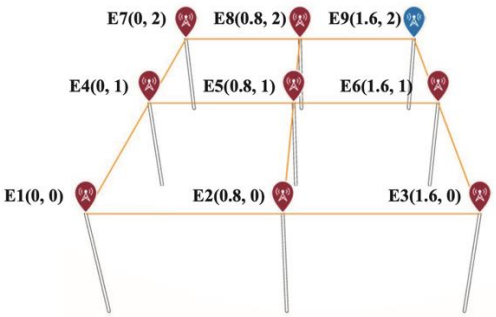

Figure 5.5: Network Topology

The carrier frequency fc is set to 60 GHz, consistent with the specifications delineated in IEEE 802.11ad and IEEE 802.11ay. Given that signal reconstruction is not an objective of our research, strict compliance with the Nyquist–Shannon sampling theorem is deemed non-essential. Consequently, we adopt a sampling frequency fs of 25 MHz, as specified in the TI sensor IWR6843 [23]. To ensure a sufficient number of sampled points with each pulse duration, we configured D_{pulse} as 200 nanoseconds. Equation (5) demonstrates our approach to achieving a relatively short pulse duration by setting $D_{symbol} = 50 \cdot D_{pulse}$, which for a 1-bit signal means that $D_{bit} = 100 \cdot D_{pulse}$. Finally, the transmission power for each node is modelled on the output capabilities of the widely used HMC 1144 power amplifier [24]. For the purposes of this experiment we statically set $V_{threshold} = 0.2mV$. In a real-world implementation, the configuration of this parameter will depend on the deployment environment. This configuration strikes a balance between technical feasibility and the practical constraints of contemporary hardware, facilitating a robust analysis within the stipulated parameters.

Latency Results: Figure 5.6 illustrates the packet-level E2E latency, which is the time it takes for a 32-bit length packet sent from the initiator to reach each node within the network through various path lengths. Due to the concurrent transmission of signals and the absence of routing, packets may follow variable paths, resulting in fluctuating E2E latency. With our current configuration, when packets of 32 bits traverse a maximum of 8 hops, the observed latency remains below 640µs. The variability in E2E latency for each network node arises primarily from uncertainties in sampling due to phase shifting, time to get to constructive superposition, and the existence of multiple transmission

paths. As in Figure 5.6, latency is rather insensitive to network diameter, increasing slowly with hop count.

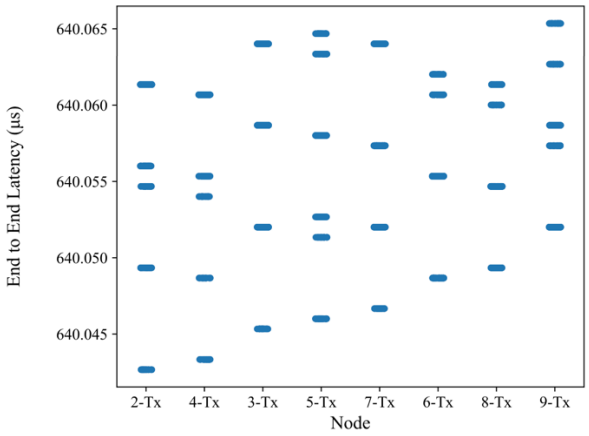

Figure 5.6: E2E latency for nodes, with fc = 60GHz, fs = 25MHz, SNR = 25dB, β = 50

Reliability Results: Figure 5.7 illustrates the Bit-Error-Rate (BER) for all nodes across varying Signal-to-Noise Ratios (SNR) with the same configuration as for the E2E latency evaluation as shown in Figure 5.6. In the evaluation of the DP modulation scheme, a compelling performance profile emerges, particularly in positive SNR environments, with the BER nearing zero at SNR levels of -10dB and above. This performance in favorable SNR conditions underscores ABL's suitability for common operational environments where positive SNR is prevalent. While there is a notable decrease in performance below -10dB, the primary focus remains on the scheme's capabilities in the typical and more frequently encountered positive SNR scenarios. This strength in standard operational contexts highlights the scheme's practical applicability and potential for robust, reliable network communication, though further study is required to understand the exact relationship between $V_{threshold}$ and ABL's performance in various SNR regimes.

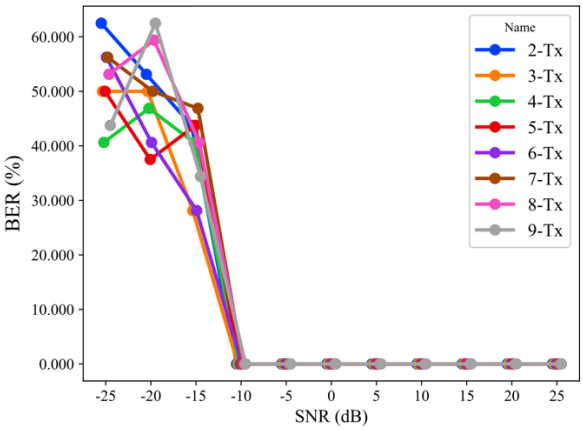

Figure 5.7: Bit Error Rate for nodes with fc = 60GHz, fs = 25MHz, SNR ∈ [−25,25] dB, β = 50

Symbol duration effect: The system's latency performance is influenced in part by its data rate. The results presented in **Figure 5.8** illustrate a decrease in symbol duration, leading to a corresponding reduction in bit duration. This change notably increases the data rate. Thereby, E2E latency is reduced, as compared to **Figure 5.6**. However, for a given pulse duration, the decrease of the symbol duration also affects reliability. As shown in **Figure 5.9**, the BER performance gets worse with low SNR configurations, because the increase of the ratio between pulse and symbol duration also increases the probability of destructive interference, which directly affects the reliability of the proposed ABL system.

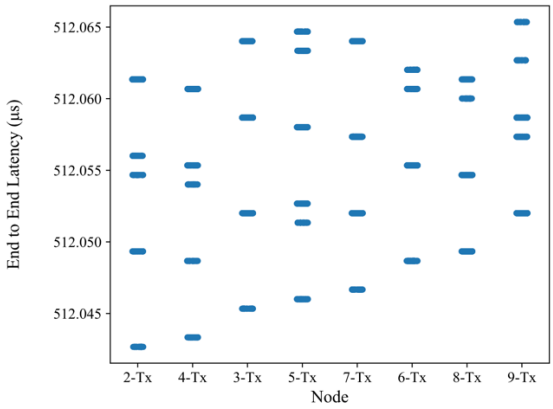

Figure 5.8: End to End latency with fc = 60GHz, fs = 25MHz, SNR = 25dB, β = 40

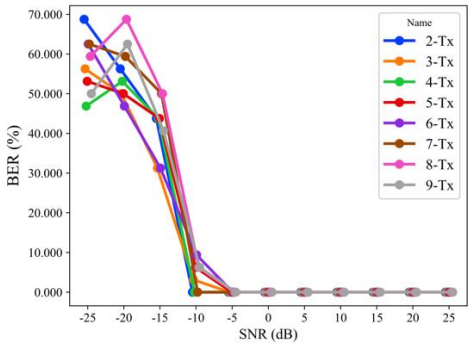

Figure 5.9: Bit Error Rate with fc = 60GHz, fs = 25MHz, SNR ∈ [−25,25]dB, β = 40

6. ISM ZeroWire Analog (2.4GHz)

This section describes a ZeroWire-like transmission design on classical sub-6GHz RF frequency bands (e.g., the 2.4 GHz ISM band). The text is based upon [3].

Inspired by the state-of-art performance of end-to-end latency of synchronous transmission, we propose a symbol-synchronous transmission design for wireless multi-hop networks. Different from the state-of-the-art solutions, we aim to develop a multi-hop protocol that provides low end-to-end latency and is suitable for long-distance communication based on radio frequency (RF) transmissions. Specifically, we propose to use widely available unlicensed RF bands (e.g., the 2.4 GHz ISM band) to broadcast information over the network. Unlike the optical band used in Zero-Wire [1], RF bands offer multiple advantages, such as a larger coverage range and better performance in non-line-of-sight conditions. Similar to Zippy and Zero-Wire [1], [17], we rely on a low-complexity OOK transmitter considering the advantage that it easily handles the interference from multiple simultaneous retransmissions of the same symbol. We evaluate its performance by means of simulation using MATLAB.

6.1 Network design

General architecture: We focus on low-latency wireless communication over a multi-hop WSN. We consider the scenario where one initiator aims to transmit data of fixed size to all other nodes in the network with high reliability and low latency. All nodes cooperatively forward the data through the network in a symbol synchronous way, allowing those further away from the source to receive it over multiple hops.

The proposed symbol-synchronous transmission way can be applied to various network topologies. To clearly explain the concept of the network and easily analyze experimental results, we simplify the network topology to a grid-like topology as shown in Figure 6.1. Figure illustrates a static grid topology where nodes are located in a

regular lattice with a fixed grid distance d. We assume there are N nodes including one initiator. All nodes use the same OOK modulation scheme and frequency band. Considering the aim of low latency, we design a new symbol synchronous transmission strategy for the considered network. Its difference from the traditional store-and-forward transmission is that the latency is not severely and negatively affected by the number of hops. Except for the initiator, every node in the network works as a transceiver. It plays two roles, which are to receive the signal and relay it. The logic of the designed strategy is that once the node receives a symbol and decodes it, it instantly broadcasts this symbol without waiting to receive the whole packet. For simplicity, **Figure 6.1** assumes that the transmission range allows only one-hop neighbors (including diagonal ones) to successfully decode symbols. As **Figure 6.1** shows, once the one-hop neighbors of the initiator, such as N_1 and N_2, successfully decode the symbol, they relay it. The node N_4 successfully receives the symbol from the node N_2, while the node N_3 receives it from concurrent transmissions of N_1 and N_2. After that, the two-hop neighbors relay the symbol thus allowing its propagation further in the network. The performance evaluation shows that with this strategy the packets can reach all the nodes with low latency, while using low transmission power.

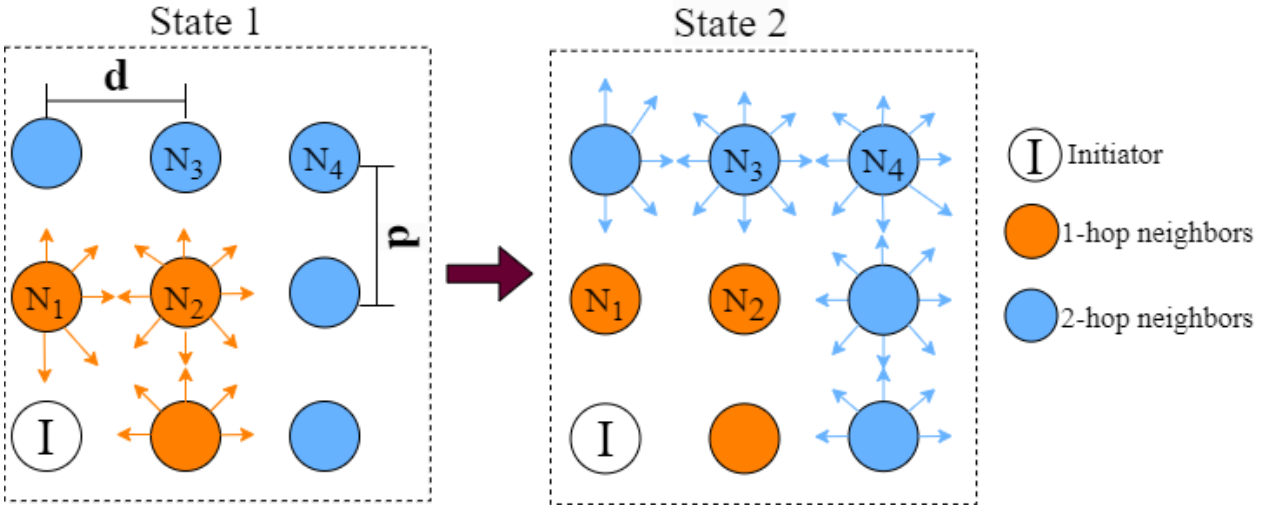

Figure 6.1: Network topology

Pulse-based OOK modulation scheme: We select pulse-based OOK as a modulation scheme. This modulation scheme is easy to implement on hardware, especially for

resource-constrained devices due to its low-complexity demodulation at low power. The inter-symbol duration is T_s. As **Figure 6.2** illustrates, when the symbol 1 is sent, the transmitter sends a short pulse with duration $T_p \ll T_s$. Otherwise, if the symbol 0 is sent, the transmitter keeps silent during T_s.

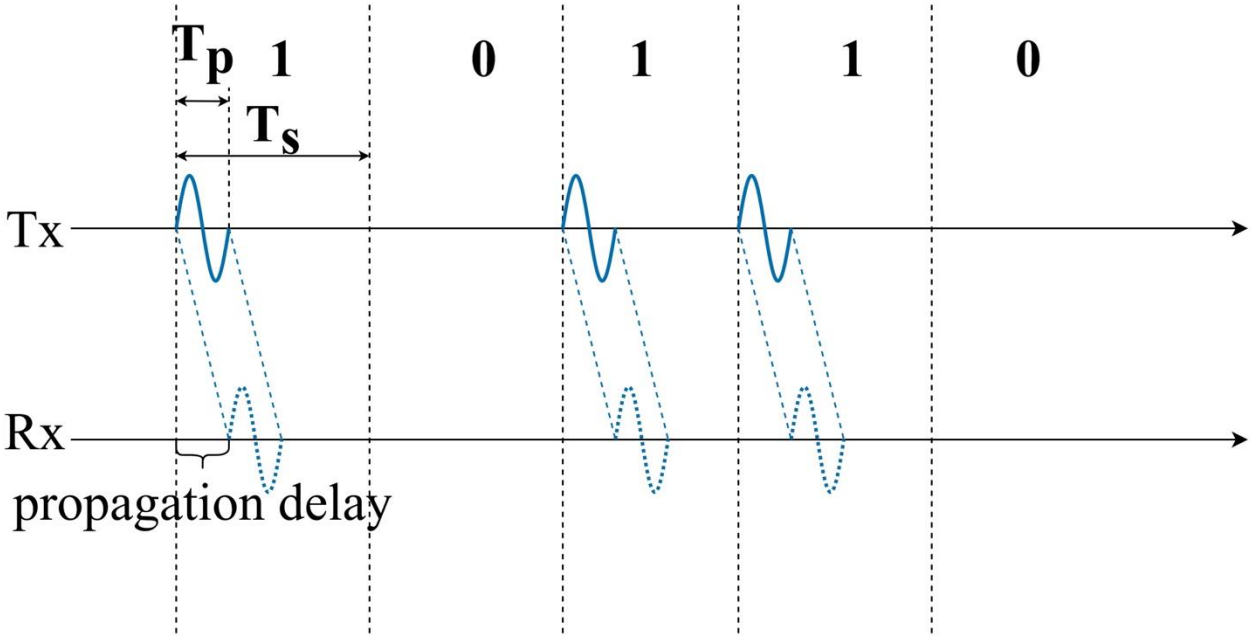

Figure 6.2: Pulse-based OOK modulation scheme

Symbol-synchronous transmission strategy: To achieve low-latency wireless communication, we design a novel symbol-synchronous transmission strategy, i.e., transmitting data symbol by symbol. Let us define the relay time r as the sum of signal propagation delay, symbol detection delay, and processing delay. We assume that the relay time is much less than the symbol duration T_s so that every node finishes retransmitting a symbol before it starts detecting the next symbol and it does not overhear retransmissions of the previous symbol while receiving the next symbol. In our analysis, we assume that the hardware processing delays are too small to impose a significant impact on the overall latency of the network, and thus are neglected. The end-to-end latency D of the network is constrained by

$$(n-1)T_s \leq D \leq (n-1)T_s + rh \quad r \ll T_s$$

where h is the number of hops from the initiator to the most distant node in the network, and n is the number of transmitted bits in a packet. Using this transmission pattern, the increase of one hop for a multi-hop network only adds r to latency.

In contrast, the end-to-end latency of store-and-forward transmission increases per additional hop by the full packet transmission time and additional waiting time related to processing and channel access at the MAC layer. The transmission logic is illustrated in Figure 6.3, if the node N_1 receives symbol 1, it relays this symbol and sends a short pulse to nodes N_2 and N_3. After a propagation delay, the nodes N_2 and N_3 respectively receive the relayed symbol at different moments caused by different distances from N_1. Then a detection procedure will be conducted to decode the received symbol. Once the node N_2 successfully decodes the signals, it will modulate the decoded symbol based on the pulse-based OOK scheme and relay the modulated signals to N_1 and N_3. The node N_3 goes through a similar procedure. After a decoding process, N_3 will relay the decoded symbol to N_1 and N_2. On the other hand, when the node N_1 detects the symbol 0, it keeps silent. Similarly, the nodes N_2 and N_3 also keep silent during the second T_s interval, corresponding to the symbol 0.

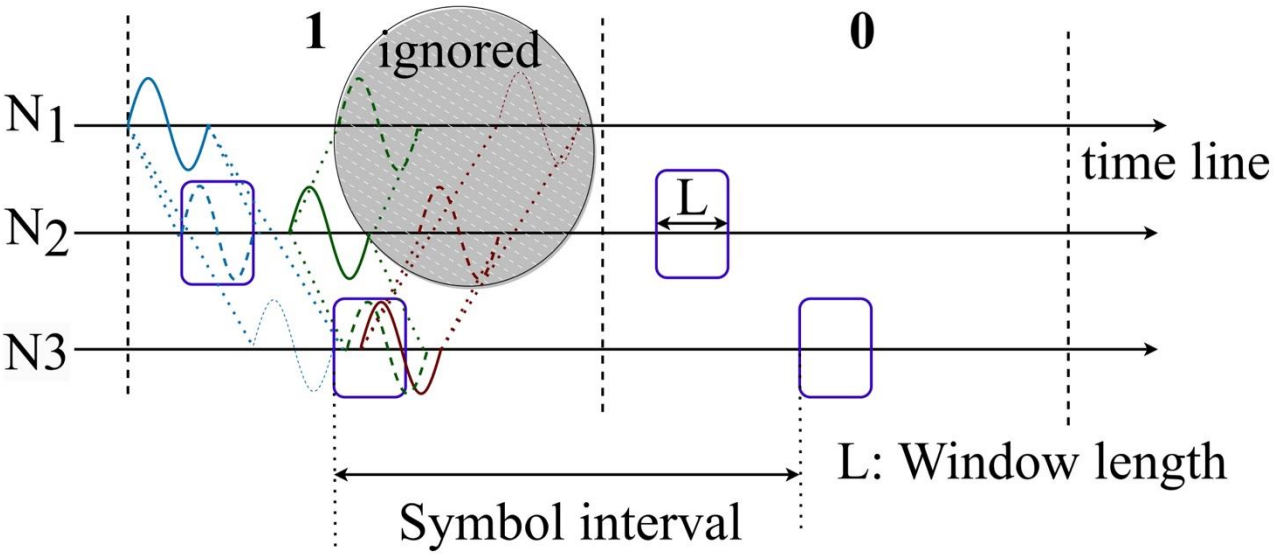

Figure 6.3: Symbol-synchronous transmission

Detector logic: This section provides some details about the detection procedures based on simulation. For the initial synchronization and detection of a new transmission,

it is assumed that each transmission starts with a bit 1 that is used as a synchronization preamble in this work. To improve robustness and avoid detecting fake symbols, a more complex preamble could be used as well. This is left for future work. As the inter-symbol duration T_s is fixed and assumed known to all nodes, a node can use the detection time of the preamble to re-synchronize its receiver. Additionally, a wake-up receiver [17] can be considered to prevent nodes from continuous listening for preamble detection. Once a receiver wakes up, it expects to receive the first symbol within a time interval of duration T_s starting from the preamble detection. Then, every next symbol will be detected in the corresponding time interval of duration T_s, which we call the symbol interval}.

In the proposed system, every node will likely receive multiple copies of each symbol representing a 1 bit, as shown in the first T_s time interval depicted in Figure 6.3. During a symbol duration, these relayed signals can be received at different moments leading to multiple detections of the same symbol. It is not only unnecessary but also a waste of energy, because a longer detection time results in higher power consumption. To address this problem, we introduce a window with length L for detection. As such, a node will only go into reception mode for a time window L at the start of each symbol interval (cf., Figure 6.3). Once the detector detects the symbol 1, it will instantly relay it and then sleep until the next symbol interval. Using this method, we can guarantee that in every time slot T_s, only one symbol is detected while its relayed copies are ignored.

Considering the signal overlapping due to various relay delays, the designed detection process is conducted several times during the time window L, improving the detection accuracy. And then we introduce a detector buffer to trigger every detection. Specifically, the received sampled signal from the transmitter is stored in a buffer until the buffer is full. After that, all stored samples in this buffer are used to make a decision by filtering the stored samples in the current buffer using a bandpass filter with a bandwidth f_{pass}, which limits the noise level. After that, the envelopes of stored samples are abstracted. Because OOK modulation is essentially an amplitude modulation, an envelope detector is an efficient way to detect amplitudes disturbed by noise. After the down-sampling of the envelope, the samples go through the amplitude comparator,

and their amplitudes are compared with a fixed threshold value determined by the receiver sensitivity. Finally, we use a voting strategy to decide which symbol is decoded. Specifically, if over half of their amplitudes are higher than the threshold value, we assume the symbol 1 is detected and the detection process is terminated until the next window time starts. Otherwise, the current buffer is emptied and starts storing the new samples for the next detection. If the detector still cannot detect a pulse at the end of the window time, we decide that the symbol 0 is detected.

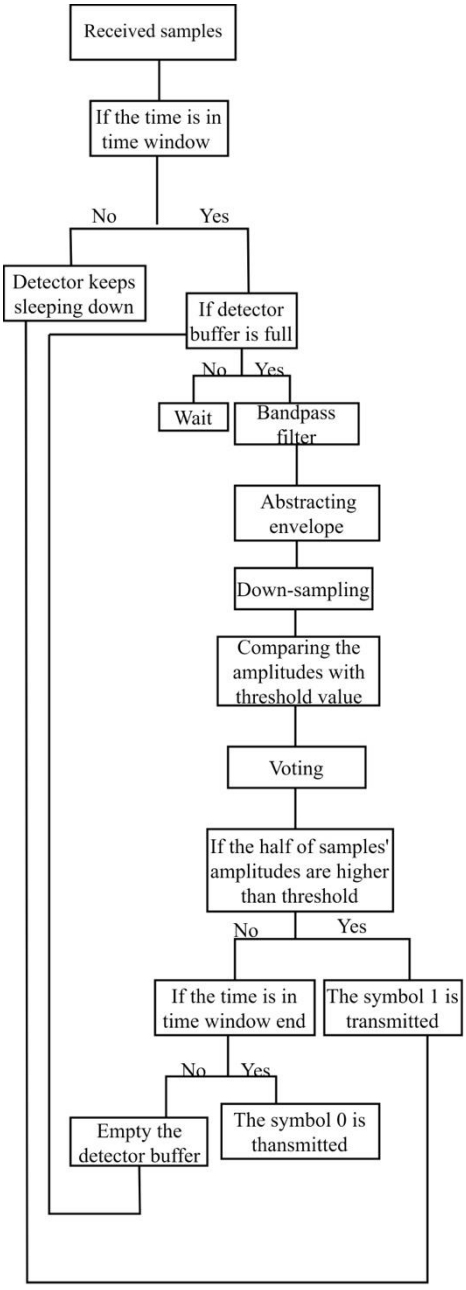

Figure 6.4: Detection scheme

6.2 System evaluation

This section describes in detail the network simulation and provides numerical results proving the network's advanced performance on end-to-end latency and reliability referring to the network BER.

Simulation setup:

We simulate a WSN in MATLAB. At the transmitter side, we simulate a carrier wave at 2.4GHz and modulate the information bits according to a pulse-based OOK scheme. The transmit power is set as 0dBm, which is reasonable for resource-constrained IoT devices. The pulse duration for a symbol 1 lasts 0.2 μs, while the symbol interval duration T_s is 10 μs. Hence, the data is sent with a bit rate of 100 kbps. In case a symbol 0 is transmitted, the transmitter keeps silent during the whole slot T_s. In addition to the modulation, we also apply a pulse-shaping filter into the transmitter to further limit the bandwidth of the transmitted signal to 10 MHz. Regarding the network topology simulation, the nodes are positioned in a square area of 4 km^2 surrounded by walls where the distance d between horizontal and vertical neighbors is varied across different experiments. In our experiments, the initiator was always located in the corner of the grid. Between all pairs of nodes, the ray-tracing channel model [25] is applied to calculate the signal attenuation and phase shift.

To simulate the analog signal transmission in MATLAB, we use a higher frequency of 9.6 GHz at the transmitter to sample the modulated signals. At the receiver side, we set the buffer size to 1000 samples to trigger every detection. A non-coherent detector is simulated consisting of a bandpass filter, envelop detector, down-sampler, amplitude comparator, and voting block. Before the amplitude comparator, a down-sampler with the frequency of 96 MHz is applied. In the simulation, the window duration time L is set as 1.875 μs to control the detector status so that in every window duration interval, the detection process is conducted 18 times.

The bandpass filter with 10 MHz is added to limit the noise level of the receiver. We simulate the receiver noise with a band-limited AWGN of the power –103 dBm corresponding to the room temperature 290K. Additionally, we set the receiver noise figure to 5 dB. The parameters' values are listed in **Table 6.1**.

Performance evaluation: This section describes and discusses the experimental results. We focus on the end-to-end latency and BER indicating network reliability. The end-to-end latency means the delay between the time when the initiator sends the first symbol of a packet and the time at which all destination nodes receive and decode the packet's last symbol. The BER is represented by the average probability of incorrectly decoding a bit across all nodes in the network. The present results do not consider retransmissions or the use of error correction codes, which could improve reliability at the cost of data rate. The integration of such strategies is left for future work. Firstly, we explore the network BER as a function of different grid distances. We place 16 nodes (4-by-4) including an initiator in a closed square area and vary the grid

Table 6.1: Parameter settings in simulation

Parameters	***Values(units)***
Data rate	100 kbps
Network area	4 km^2
Sampling frequency of transmitter	9.6 GHz
Carrier frequency	2.4 GHz
Pulse duration	0.2 μs
Transmitter power	0 dBm
Noise power	-103 dBm
Noise figure	5dB
Receiver sensitivity	-90dBm

Window time	1.875 μs
Buffer size	1000 samples
Number of times for detection	18
Sampling frequency of receiver	96MHz
Bandwidth	10MHz

distance from 50 m to 200 m. We randomly generate and send 100 packets consisting of 64 bits each. The results are depicted in **Figure 6.5**. As expected, BER increases with the grid distance. As the grid distance i.e., the distance between horizontal and vertical neighbors in the grid, increases, propagation loss will increase, due to longer links, but also the number of hops to reach further nodes will increase. We can see the network BER is lower than 1% when the grid distance is below 100 m. Particularly, when the grid distance is 50 m, we cannot see any errors for 100 packets. On the grid distance of 60m, the BER is 0.04%. When the grid distance is over 100 m, the network BER increases sharply, because the transmission power of 0 dBm is too low to support decoding with high accuracy even for the nearest nodes from the initiator (i.e., one-hop nodes).

In addition, we evaluate our network for end-to-end latency, for packet sizes equal {64, 128, 256, 512} bits across the 16-node network with grid distance ranging from 50 m to 200 m. As a function of packet size, the latency almost linearly increases as shown in Figure. According to **Figure 6.6,** our network end-to-end latency is in the order of a few milliseconds. Specifically, the average latency to transmit a 512-bit packet is around 5 milliseconds. For a small packet of 64 bits, the average latency is only around 0.63 milliseconds. Moreover, the latency does not seriously increase by grid distance increase. The latency increase caused by distance increase from 100 m to 200 m is of orders of a few microseconds as shown in **Figure 6.6**. The slight increase is due to the increased propagation delay and is negligible.

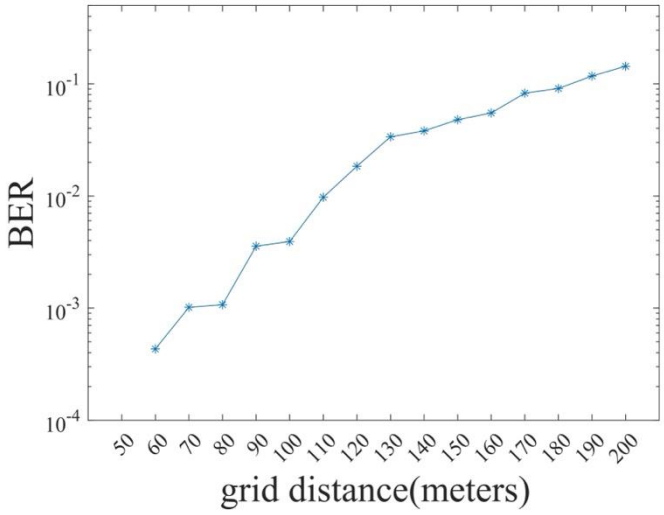

Figure 6.5: Average BER for various grid distances

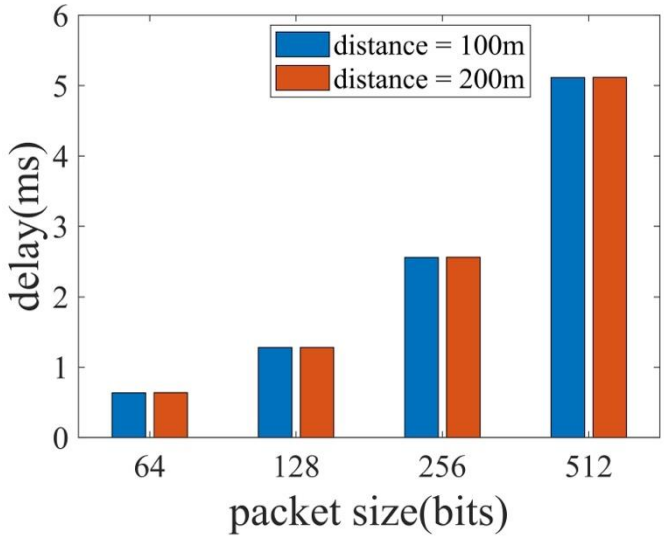

Figure 6.6: Network latency for grid distance 100 m and 200 m

To explore the latency behavior as a function of the number of hops, which is the main advantage of symbol synchronous transmission compared to store-and-forward, we perform another experiment to expand our network into 64 nodes (8-by-8), using the same network topology as shown in **Figure 6.1** with a grid distance of 100 m. **Figure 6.7** depicts the latency of each node depending on their distance from the initiator in terms of the number of hops. While the number of hops increases as a function of the distance, the figure shows that the increase is limited to about 2

μs per hop (i.e., the pulse duration of 0.2 μs). This results in a latency increase of about ~0.3% per hop. This is an immense improvement compared to store and forward networks, which add at least 100% latency for each additional hop.

After exploring the distance effect on network BER, we are also concerned about the effect of the number of nodes on network reliability in terms of BER. In this experiment, we fix the grid distance and change the number of nodes instead. Specifically, we first fix the grid distance to 50 m and transmit 100 64-bit packets through the network consisting of {16, 25, 36, 49, 64, 81} nodes. Subsequently, we fix the grid distance to {75, 100, 125} m, and repeat the above experiment. The results are illustrated in **Figure 6.8**. According to **Figure 6.8**, we can conclude that the increase in the number of nodes harms network reliability. This is due to the fact that a more dense deployment results in more nodes relaying the same symbol simultaneously, which can cause destructive interference. We leave addressing this problem for future work. Despite these negative effects, our network can still achieve reasonable BER below 5% when the number of nodes is less than 50 and the grid distance is less than 75 m.

Performance comparison: In this section, we compare our network performance with other state-of-the-art low-latency networks. In **[1]**, performance evaluation shows that Zero-Wire can achieve 99% reliability for 32-bit frame reception. In **[17]**, Zippy reaches 94.6% for 16-bit frame reception. According to our results, the proposed design with a grid distance of 60 m can achieve BER of 0.04%, which results in $(1 - 0.0004)^{32} \approx 98.7\%$ frame delivery ratio for a 32-bit frame based on an assumption of statistical independence of subsequent demodulation attempts. Although the reliability performance is similar, the data rate for the proposed design is significantly higher, 100 kbps in comparison to 20 kbps supported by Zero-Wire and 1.364 kbps for Zippy. Besides, a 2.4 GHz frequency band provides better coverage and propagation properties, allowing reliable data transmission over larger distances.

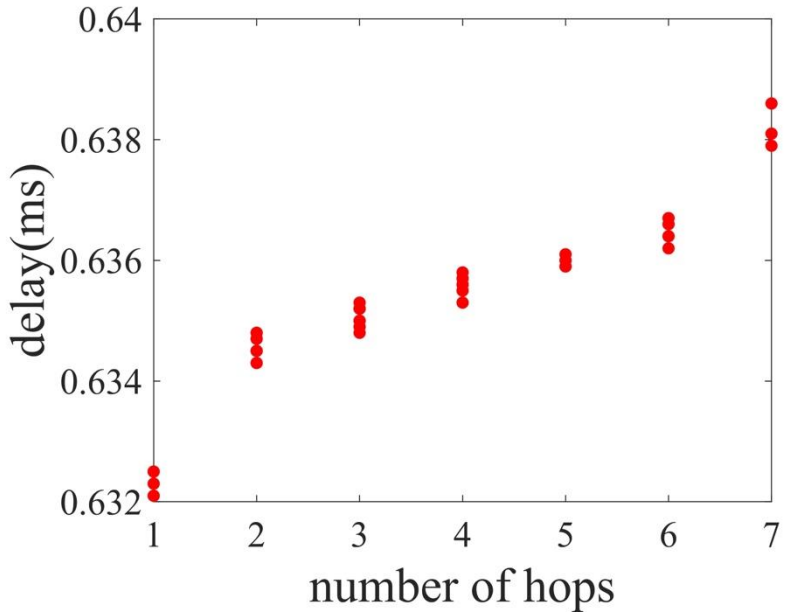

Figure 6.7: Effect of the number of hops on latency for 64 bits and grid distance 100m

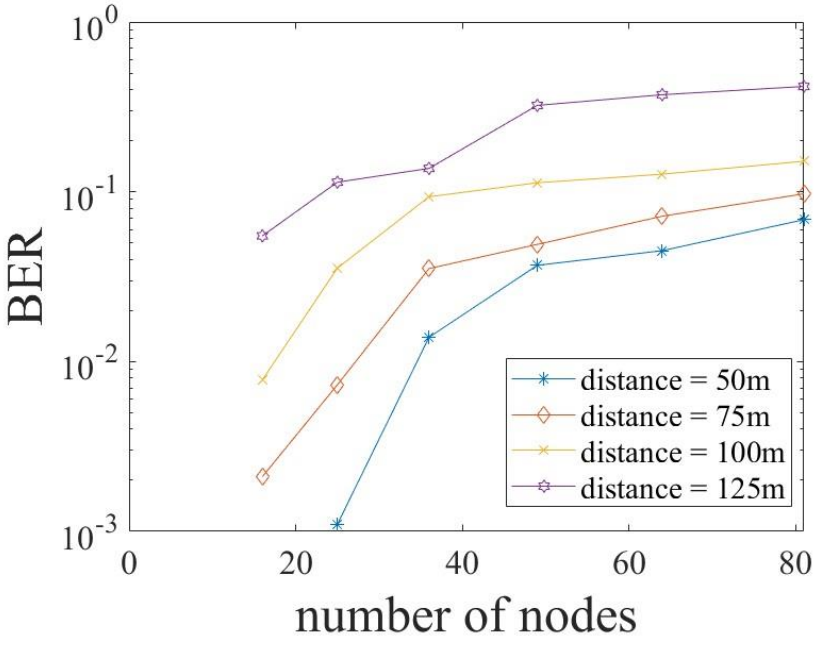

Figure 6.8: Average BER for various numbers of nodes

However, it should be noted that in contrast to [1], [17], our solution is only tested using simulation and lacks convincible performance evaluation based on real

hardware. However, our experiment results demonstrate for the first time how the order of magnitude improvements are possible by properly designed hardware since traditionally empirical platforms face considerable incidental limitations in their implementation. In our future work, we will conduct such experiments using a software-defined radio (SDR) for a fairer comparison.

7. VHF ZeroWire Analog (40MHz)

Aside from modulation-specific properties, the root of the energy-latency trade-offs found in conventional low-power mesh networks is the need to reduce power consumption by duty cycling the radio receiver using time-synchronized wake-ups or preamble sampling. This section describes a VHF platform called CaIN that eliminates the need for such Radio Duty Cycling (RDC) by reducing idle listening power. This approach is hence not a direct translation of the ZeroWire protocol, taking a different approach towards transceiver re-design that drives down latency in low-power mesh networks. The text in this section is based on [4].

Current WMN radios alternate between low-power states in which they cannot receive data and more power-hungry states in which they can. Such duty cycling is necessary to reconcile their power consumption in receive mode (several mW) with the average energy budget available over multiple years (tens of µW). As data can only be communicated if a transmitter and a receiver agree on when to do so, WMNs coordinate rendezvous points, for example by synchronizing wake-ups in time [6, 7], or by preamble sampling [26, 27]. Such duty-cycled rendezvous increases end-to-end latency for aperiodic or event-driven traffic, which often scales to many seconds for large mesh networks. This precludes the use of these technologies in time-sensitive applications. The core scientific contribution of this workis to introduce the paradigm of near field mesh networking, a novel approach to low power and low latency networking. While the near field is limited in range to a few wavelengths due to its reliance on capacitive and inductive effects, those effects are dramatically stronger than the far field. This higher received signal strength simplifies transceiver design, enabling sub-

microwatt power consumption in receive mode. Near-field communication has a long track record in the form of Radio Frequency ID (RFID) [28], though the range of RFID is limited to a few meters even with powerful transmitters and directional antennas. CaIN addresses this range limitation with a two-fold approach; firstly by operating at VHF frequencies to maximize near-field range and secondly building a WMN to enable the coverage of larger areas.

The proposed paradigm is supported by two engineering contributions; (i.) an open source hardware and software reference design, that enables further experimentation with near-field mesh networking and (ii.) a systematic evaluation of the approach in real-world conditions. The current prototype of CaIN operates between 20MHz and 80MHz. This carrier wave is modulated using On/Off Keying (OOK) and Manchester coded. The frequency of the carrier wave is configurable in software and may be tailored to suit the specific environment and available frequency bands.

CaIN offers a unique and complementary performance profile in comparison to Radio Frequency ID (RFID), 802.15.4 and Bluetooth Low Energy (BLE). Its ultra low receiver power enables always-on operation, thereby eliminating the latency that arises in mesh networks due to receiver duty cycling. Furthermore, reasonable transmitter power consumption together with omni-directional antennas enable the construction of a multi-hop wireless mesh with dramatically lower end-to-end latency than previous approaches.

CaIN is a hardware and software platform reference design of near field mesh network running at 40 MHz. The platform transmits information using strong Capacitive and INductive effects that can be detected by a passive receiver front end, thereby side-stepping the need for intermittent receiver usage. Given that near field effects are limited to a few wavelengths, we implement our proof-of-principle transceiver in the VHF frequency band, achieving a unique performance profile including: 20kbps link layer throughput, sub-µW receiver power consumption, sub-ms wireless wake-up and 17m per-hop range. Building upon these properties, we realize a novel mesh network with a worst-case latency of under 40ms for a 3 hop network. To promote replication

and further work, we have made the hardware and software of our reference implementation open source.

7.1 Background

Radio Duty Cycling: Conventional radios use significant power when listening for incoming transmissions and must therefore be duty cycled in order to maintain an acceptable power profile. Two archetypal approaches have emerged to managing radio duty cycles and rendezvous: preamble sampling and time synchronization.

With preamble sampling [27], one of a pair of devices announces its intent to communicate data by transmitting a bit pattern of length L in time, while the other device samples the channel at a rate of R, where R ≤ L. Preamble sampling enables the receiver to be duty cycled at the expense of additional power consumption and latency for the sender. Longer preambles will result in lower power consumption for receivers at the expense of latency. Similarly, low-power probing [20] approaches have receivers announce their willingness to receive data, thus also reducing spurious listening at the expense of latency.

Time Synchronized Network (TSN) protocols such as 6TiSCH [6] schedule per-node transmit and receive slots and thereby eliminate the need for preamble sampling and further reduce power consumption. However, latency is increased for unpredictable traffic flows, as transmitters must wait for a scheduled transmit slot before they are able to send. In the case of multi-hop mesh networks, this latency scales with the diameter of the network, often to several seconds.

Wake-up Radios are a class of always-on, ultra low power radio design. As such radios are considerably less sensitive than their more power consuming counterparts, they are typically used to communicate a special-purpose wake-up signal that initiates subsequent communication using a conventional primary radio [29]. For the same reason, WuRs often depend on relatively powerful transmitters [29]. Since a WuR is always on, latency can be reduced in comparison to the RDC approaches. This deliverable demonstrates that WuR-like radio design, when combined with VHF

communication and near-field meshed deployments, enables considerable improvements to the latency-energy trade-off encountered in low-power networks.

WuRs operating in the µW range share many common design elements, being based around envelope detectors, passive voltage multiplication and digital comparators to demodulate an On/Off Keyed (OOK) wake-up signal. For instance, Bdiri et al. [30] introduce a 868MHz receiver, which achieves a range of 2.5m, data rate of 1.2kbps and a power consumption of 690nW. A CC11011 transmitter is used to send the wake-up signal consuming 74mW. Marinkovic et al. [31] introduce a 433MHz WuR, which achieves a range of 10m, a data rate of 5.5kbps and a power consumption of 270nW. An ADI ADF70202 was used to provide a wake-up signal, consuming 85mW. Magno et al. [32] introduce an 868MHz WuR which achieves a range of up to 50m, a data rate of 10kbps and receiver power of 1.2µW. A TI CC430F614x3 provides the wake-up signal, consuming 105mW. For a complete review of wake-up radios, we refer the reader to Piyare et al. [29].

Radio Frequency ID (RFID): Radio Frequency ID (RFID) is designed to support simple and low cost tags that may be read and written at short range by a powerful reader. In the case of passive tags (i.e. battery free), Wireless Power Transfer (WPT) from the receiver is used to provide all operational power. RFID is available in a range of frequencies: LF (135 kHz), HF (13.5MHz) and UHF (868/915MHz). Near Field Communication (NFC) is a strict subset of the HF RFID variant [28]. LF and HF tags depend upon the strong coupling effects that occur within the near field and offer ranges between a few CM and a few meters, while UHF tags, for example, the WISP platform depend upon the far field effect and offer operational rages of around 3 meters using a conventional reader. RFID tags communicate using backscatter, modifying the properties of their antenna to selectively reflect the incident signal from the receiver and thereby transmit a response to the reader at orders of magnitude less power than an active radio. While RFID offers a simple and low power solution to communication, it is inherently short range and due to the requirement for a powerful reader, typically operating at 10s of dBm and consuming multiple Watts, RFID is limited to star topologies and cannot practically be meshed.

Requirements: The overall goal of CaIN is to enable the construction of low latency mesh networks at ultra-low power. This leads to three technical requirements:

1. Ultra Low Power Receiver: RDC inevitably increases per-hop latency. An ultra-low power always-on receiver is required which does not add significantly to the deep sleep current of a conventional MCU (order of 1µW).

2. Integrated Transceiver Design: The low data rate and high transmitter power of conventional WuR requires the use of a separate radio for wake-up and communication, increasing size and design complexity. An integrated transceiver is required which supports both sub-µW listening and similar end-to-end data rates to conventional radios (order of 10kbps).3. Transmitter Range must also be maximized to preserve freedom in deploying and designing mesh networks. This also precludes the use of directional antennas as is common in RFID. Specifically, we aim for an omnidirectional range of at least several meters.

In the following Section we elaborate on the design of CaIN, a system which fulfils these requirements.

7.2 Design

In this section, we describe the design of CaIN. We first describe the physical layer design of CaIN, followed by the CaIN data link layer, Together these two layers are sufficient to realize a simple single-controller wireless bus. The higher layers of the protocol stack, from Network to Transport are the subject of future work.

Physical Layer: The physical layer design of CaIN is illustrated in the block diagram of Figure 7.1.

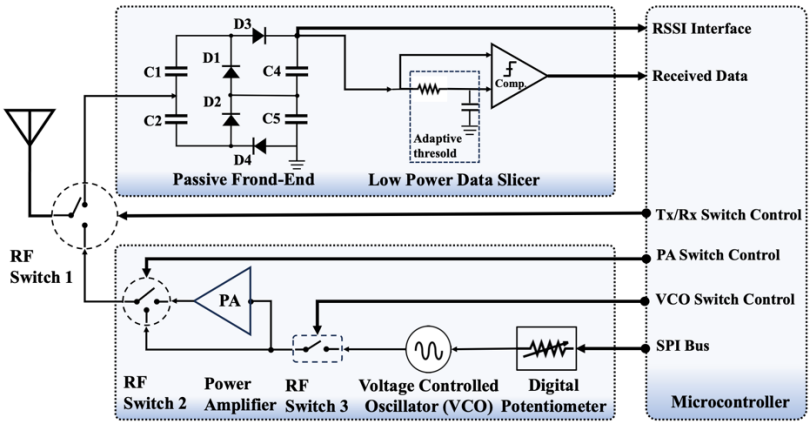

Figure 7.1: High-level block diagram and of a CaIN transceiver.

CaIN takes a very simple approach to transmitting data, which in turn enables the creation of a nearly-passive receiver. A Voltage Controlled Oscillator (VCO) generates a Very High Frequency (VHF) carrier wave in the range of 20 to 80MHz. This enables access to three unlicensed frequency bands at 27, 35 and 40MHz4. Certification in one of these frequency bands will be the subject of our future work. The frequency of the VCO is selected by the application MCU via a digital potentiometer connected to the Serial Peripheral Interface (SPI) bus. The application MCU controls switching between receive and transmit modes via RF Switch 1. The output of the VCO may be amplified by a Power Amplifier (PA) which is controlled by RF switch 2. Data is modulated by On/Off Keying (OOK) the VCO output using a RF Switch 3. Experimentation has shown that data rates of up to 4.9kbps are achievable without amplification and 19.8kbps with the amplifier active.

The CaIN front-end is based on a Greinacher voltage multiplier, which rectifies the received signal and increases its voltage to a usable level. This is a common design in passive RF front ends such as those found in WuR [29] and in Wireless Power Transfer (WPT) systems [55]. By building our transceiver in the VHF frequency range, we side-step much of the optimization complexity that is typical in passive and ultra low power receivers in the UHF and ISM* bands (868MHz and 2.4GHz). At ranges of up to a few wavelengths (7-30m for the CaIN transmitter), near field coupling effects result in a strong received signal strength even at modest transmit power. The Greinacher voltage

multiplier operates as a capacitive charge pump circuit. The optimal capacitance value can be calculated using the following formula:

$$C = \frac{I \cdot T}{\Delta V}$$

Where C is the size of the capacitors in the voltage multiplier, I is the load current, T is the period of the received waveform and ΔV is the minimal voltage swing that is necessary to overcome the internal hysteresis of the low-power data slicer. In practice, capacitors may be selected slightly above the calculated values to provide a buffer against suboptimal conditions. As can be seen from Equation 1, high carrier frequencies (supporting fast symbol rates) demand low capacitance, whereas lower carrier frequencies (offering extended range) necessitate higher capacitance. In the current prototype of CaIN, these components are manually tailored to match the carrier frequency.

Where Low-power data slicing: Conventional radios use high-speed (and therefore high-power) Analogue-to-Digital Converters (ADCs) to demodulate the received signal. However, the low power data slicing module of CaIN takes a different approach, using an sub-µW comparator to threshold the analogue signal and thereby convert it into a series of symbols. An RC low pass filter is applied to one of the comparator's inputs to support dynamic thresholding of the input signal and increase immunity against background noise. As with the passive front end, the values of R1 and C5 are tailored to suit CaIN's symbol rate as follows:

$$f_c = \frac{1}{2\pi RC}.$$

fc is the filter's cutoff frequency, which we set to 330 Hz, well above common sources of AC interference, but well below CaIN's minimal data rate. The output of the data slicer is provided to the CaIN data link layer running on the application MCU. In the current

version of CaIN, the values of analog components in the passive front-end and data slicer are fixed. However, dynamically tuning them forms part of our future work.

Data Link Layer: The CaIN data link layer implements a simple single controller addressable bus structure. That is, a single controller device initiates all traffic by transmitting a request to some other device. That request is then routed through the network across multiple hops until it reaches it destination. Next, the destination node generates a reply packet, which is then routed in the opposite direction, i.e. towards the controller node.

CaIN encodes binary data using Differential Manchester coding, wherein each bit is represented by a transition or a lack thereof in the signal. This approach provides synchronization, noise immunity and a degree of error resilience. The self-clocking character of Manchester coding is a good fit with CaIN, wherein the receiver must decode data encoded at different data rates (from 4.9 to 19.8kbps).

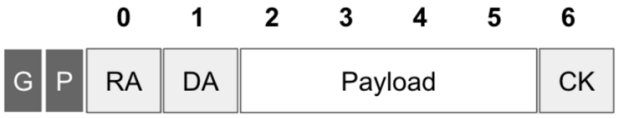

Figure 7.2: CaIN Packet Structure

The data link layer frame format of CaIN is shown in **Figure 7.2**, wherein:

- G is 0.8ms (4b) of dead-time used to delineate packets.
- P is 0.8ms (4b) of preamble to wake the receiver.
- RA is an 8b Receiver Address for the next hop.
- DA is an 8b final Destination Address.
- Payload is 32b of freeform application data.
- CRC is an 8b checksum to detect corrupt packets.

This small packet structure is heavily optimized for low latency operation, enabling CaIN to maximally exploit latency gains arising from its always-on receiver. As we implement a single-controller bus, no source address is required.

The forwarding behaviour of CaIN is simple. At each hop, the receiver RA will first validate checksum CK, discarding corrupted packets. If the packet is intact, a lookup is performed on RAs routing table for the destination address DA. If DA is known, a new packet is then generated and retransmitted to the next hop on the route to DA. As the current version of CaIN implements a unicast single-controller bus, there is no source of contention, though interference may occur from co-located RF sources.

CaIN does not tackle the problem of finding and dissem- inating routes accross the network. For the purposes of the mesh experiments performed in this deliverable, multi-hop routes are known a-priori. In our future work we will investigate how existing approaches such as the Routing Protocol for Low-Power and Lossy Networks (RPL) [33], 6TiSCH [6] and Ad-Hoc On Demand Distance Vector (AODV) [34] can be applied for static and mobile scenarios respectively.

7.3 Implementation

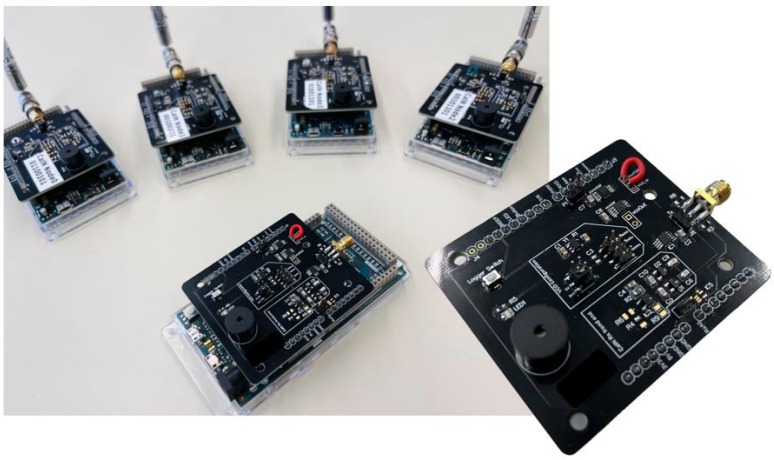

Figure 7.3: Physical Prototype of CaIN in Arduino Shield Layout

This section describes a prototypical implementation of CaIN, relying on a embedded microcontroller to control the low-power radio front-end under test.

Software: The CaIN transceiver has minimal requirements of the Application MCU. In terms of IO, CaIN requires an SPI connection to the Transmitter and three General Purpose IO (GPIO) pins to switch the transceiver mode, send and receive data. The application MCU also hosts the CaIN data link layer.

Following initialization, the MCU uses the RF switch that connects the antenna to the receiver front-end to place the transceiver into receive mode. the MCU may then attach an interrupt handler to the RX pin and enter sleep mode, ready to be activated by a received CaIN message. Once awake in receive mode, the MCU will access the digital symbol stream from the data slicer to decode incoming packets. As the receiver is always on, yet consuming very little power, the need for RDC is eliminated.

When acting as a transmitter, the transceiver's MCU activates the RF switch via GPIO and connects the antenna to the transmitter's front end. The MCU fine-tunes the oscillator frequency by adjusting a digital potentiometer. Data may then be modulated by using the TX pin to toggle the RF switch, thereby modulating the carrier frequency for signal transmission. Transmit power and maximum symbol rate can be increased by activating the embedded amplifier (from 4.9 to 19.8kbps). We believe that efficiency and simplicity can be further improved by moving the full implementation of CaIN to a dedicated transceiver module offering a packet-based interface. This will be a subject of our future work.

Hardware: The current version of CaIN is shown in **Figure 7.3** and matches the footprint of an Arduino shield enabling easy prototyping with a wide range of application MCUs. The hardware Bill of Materials for the CaIN transceiver along with the component cost in units of 10K is shown in **Table 7.1** below. This excludes both the application MCU and antenna, as we expect these to be customised based on the needs of the application.

Table 7.1: CaIN BoM in Units of 10K

CaIN Prototype Components	Value	Price @10K(USD)
Rectifying Diode	SMS7630- 005LF	0.57
RF SPDT Switches	ADG918	2.46
RF SPST Switch	ADG902	1.49

Analog Comparator	TLV7031	0.17
VCO	LTC6905	1.15
RF Amplifier	TRF37D73	0.43
Resistors and Capacitors	NA	0.23
Antenna Connector	SMA- EDGE-S	1.59
	Total:	8.09

As can be seen from Table 7.1, even when fabricated from discrete components, the cost of a CaIN transceiver is low. As no exotic components are required, it should be possible to realise a future design of CaIN as a miniaturised System on Module at a significantly lower price point.

7.4 Evaluation

Transceiver characteristics: The performance characteristics of the CaIN transceiver are shown in Table 7.2. CaIN listens for incoming messages with a power consumption of 558 nW. On receiving a message, the application MCU wakes from sleep mode in under 0.8 ms. In this mode it is capable of receiving data at rates between 4.9kbps and 19.8kbps. At 4.9kbps packet reception takes 6.8nJ, while at 19.8kbps, it takes only 1.67nJ, several orders of magnitude lower than a conventional radio. The transmitter has two power settings, 23.5mW which supports a data rate of 4.9kbps and consumes a total of 286µJ per transmission and 201.7mW, which supports a data rate of 19.8kbps and consumes a total of 609µJ per transmission; roughly equivalent to conventional PAN and LPWAN radios respectively.

Table 7.2: CaIN Transceiver Performance Characteristics

Metric	Unamplified Transmitter	Amplified Transmitter
Throughput Avg. packet speed	 4.92 Kbps	 19.8 Kbps
Timings Preamble latency Packet transmit Total On-air time	 0.80 ms 11.39 ms 12.19 ms	 0.20 ms 2.82 ms 3.02 ms
Power Draw RX Passive	 558 nW	 558 nW

TX power	23.5 mW	201.7 mW
Energy RX Packet TX Packet	 6.802 nJ 286 µJ	 1.685 nJ 609 µJ

Table 7.3 provides detailed timing characteristics for the CaIN transceiver in its 4.9kbps and 19.8kbps modes. As can be seen from the table, CaIN enables wireless wake-up in under 0.8ms and address matching in as little as 1.61ms. This enables the application MCU to remain in deep sleep mode while being responsive to incoming messages. Furthermore, as address matching can be performed in under 2ms, nodes can return to deep sleep mode quickly in cases where they are awakened by a message that was not intended for them.

Table 7.3: CaIN Transceiver Timings

CaIN Mode	**@ 5 Kbps**	**@ 20 Kbps**
Preamble Delay	0.8 ms	0.2 ms
Address Match	1.61 ms	0.40 ms
Payload Send	9.78 ms	2.42 ms
Switching Time	10.60 µs	11.81 µs
Total	12.19 ms	3.03 ms

Pairwise Performance: Figure 7.4 shows the strength of the received signal in mV along with the noise floor for a receiver at distances of 1 to 15m from the transmitter with carrier frequencies of 20, 40 and 80MHz. All experiments were conducted outdoors in a suburban setting with a transceiver height of 2m using a 1.2m telescopic monopole antenna. Each data point was recorded 10 times and the results were averaged.

The figure embodies a trade-off between the range of the near field, which decreases linearly with frequency and the sub-optimality of our 1.2m antenna, which performs increasingly poorly at lower frequencies. For this hardware configuration, 40MHz provides the most attractive trade-off between range and received signal strength, remaining well above the noise floor at a distance of 15m. We therefore use a 40MHz

carrier for the remainder of this evaluation. Achievable range also depends upon required reliability and transmitter power, as will be explored in the following section.

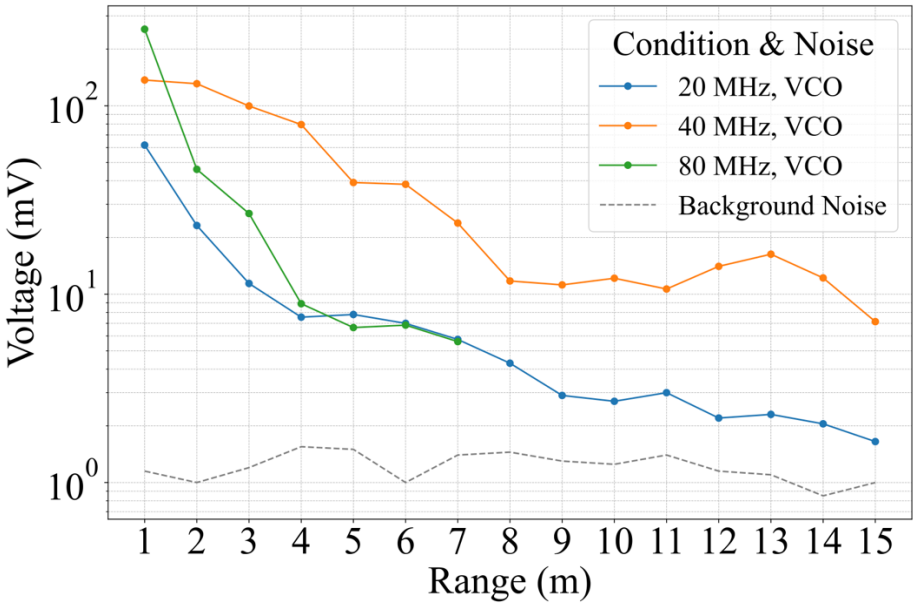

Figure 7.4: Received signal strength for a pair of transmitters at 20, 40 and 80MHz and distances of 1 to 15m.

Reliability: Figure 7.5 shows CaIN's Packet Delivery Ratio (i.e. the % of successfully received packets) for increasing packet sizes from 4 to 8B with increasing range between the transmitter and receiver. The experiments are executed twice using the two power settings and speeds of the CaIN receiver 201.7mW / 19.8kbps and 23.5mW / 4.8kbps.

As can be seen from the figure, PDR at first falls gradually with distance (until 10m at 23mW and 17m at 202mW), after which it drops precipitously to zero. This is as expected for near-field communication which is inherently limited in range. Naturally, small packets result in the highest PDR as there is less opportunity for bit-flips to occur.

Until the limit of near field communication, there is a trade-off between range and reliability. In line with current radios, we aim for a base-level of 75% PDR, resulting in an effective range of 7m at 23mW and 17m at 202mW.

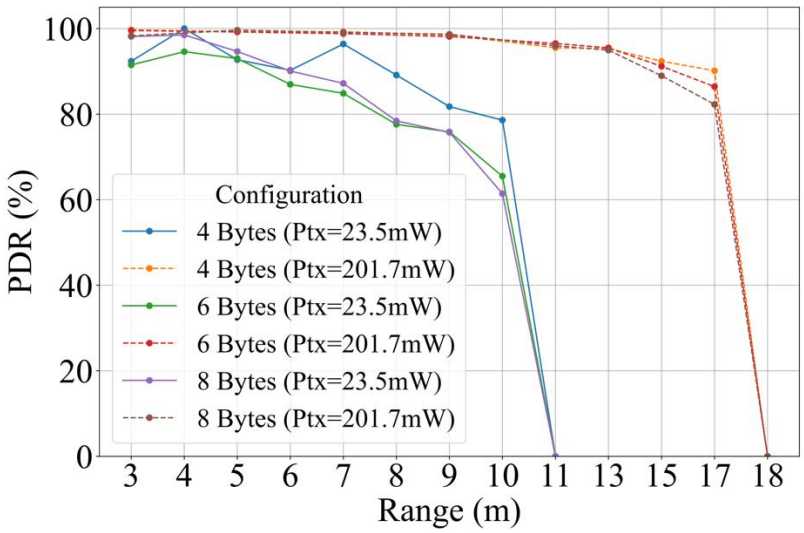

Figure 7.5: Packet Delivery Ratio (PDR) for 4, 6 and 8B at a range of 1 to 18m.

Impact of Density: For near-field communication, the number of receivers near the transmitter affects signal strength due to coupling. For every extra node that is placed between a transmitter and receiver, a notable drop in received signal strength occurs. In conventional wireless mesh networks, which rely on radiative power transfer rather than near-field coupling, such effects of deployment density on signal strength do not occur. This relationship is quantified in **Figure 7.6** wherein we analyze (i.) a receiver that is part of an increasing number of nodes positioned at even angles along a circular arc corresponding to a 2-metre transmission range, and (ii.) a receiver that is part of a line of nodes spaced 2 meters apart, thus studying how signal strength evolves as deployment density increases.

As can be seen from **Figure 7.6**, signal strength is a decreasing function of deployment density surrounding the transmitter, but the maximum effect is 18% as the deployment evolves from a minimally to a maximally dense scenario, and that it decreases by no more than a few millivolts with every additional node. Interpreted in the context of **Figure 7.4**, the impact of the maximally dense deployment scenario on signal strength hence is dramatically smaller than that of positioning the transmitter-receiver pair merely one more meter apart. In mesh deployments, this adverse effect would easily be offset by the reduction in transmission distance that occurs due to a denser deployment.

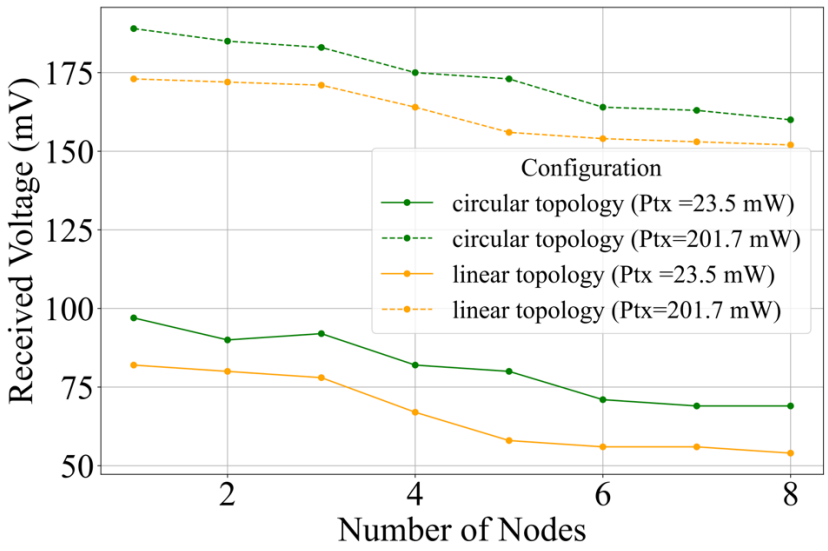

Figure 7.6: Network densification affects signal strength due to near-field coupling, though impacts are manageable.

Mesh Latency: Figure 7.7, evaluates the latency of a network of two to four nodes (i.e. one to three hops) deployed in a line topology with a spacing of 2m between each node. 600 packets were transmitted for each data point.

As the single controller bus implemented by CaIN has no contention, the observed results are deterministic, matching the timings provided in Table 7.3 and scaling linearly with hop-count as expected. It is important to note however that, unlike conventional networks such as SmartMeshIP or Bluetooth Low Energy 5, the CaIN receiver is always on and does not need to be duty-cycled. In the following section we compare the performance of CaIN to representative contemporary radios under different duty cycle settings.

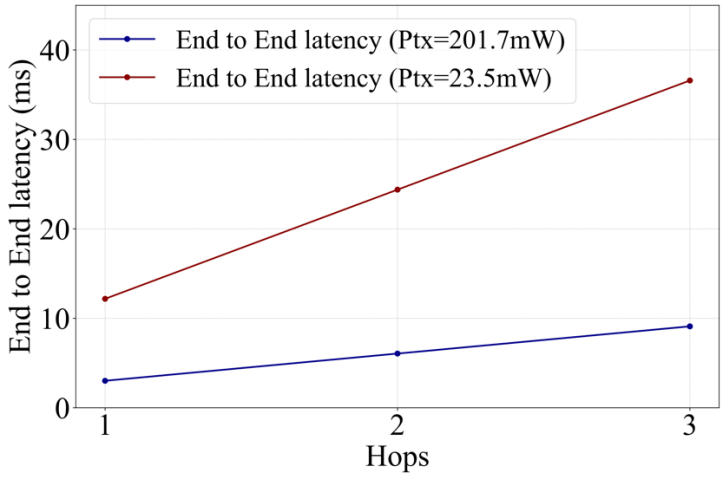

Figure 7.7: End-to-end Latency for networks of 1 to 3 hops.

Comparison to contemporary radios: In this section, we use the measured performance profile of CaIN together with manufacturer provided data for two state-of-the-art low power wireless network chips.

SmartMesh IP (SMIP): The LTC5800-IPM from Analog Devices6, is an 802.15.4 transceiver that implements TSMP [7], a 6TiSCH-like [6] time-synchronized mesh network. This radio has a base bandwidth of 250kbps, 7.25ms TX/RX slot length and supports packets of up to 128B. It consumes an average of 1.6mW during an active receive slot and 1.96mW during a transmit slot. All performance numbers were drawn from the data sheet.

Bluetooth Low Energy (BLE): The nRF52840 from Nordic Semiconductor is a BLE transceiver. We used a base bandwidth of 1mbps, 3.33ms TX/RX slot lengths, and packets of 32B. This radio consumes an average of 1.2mW during both TX and RX slots. Performance data was drawn from the Nordic BLE power estimator tool.

Figure 7.8 plots the lower bound on latency for BLE, SMIP and CaIN. For BLE and SMIP this requires a 100% duty cycle, while the CaIN receiver is always-on throughout the experiments. As can be seen from the figure, CaIN is competitive with these protocols in terms of latency even at 100% duty cycle offering a latency of 3.03ms at 19.8kbps and

12.19ms at 4.9kbps. However, this arises in large part due to CaIN's smaller 7B packet size, versus 32B for BLE and 128B for SMIP.

The key difference between CaIN and contemporary low power wireless radios becomes apparent in **Figure 7.9** which shows the power consumption for all technologies at 100% duty cycle. As can be seen from the figure, the power consumption of SMIP and BLE dwarfs that of CaIN by several orders of magnitude.

Considering **Figures 7.8** and **7.9** together, it becomes apparent that CaIN offers an excellent power/latency trade-off irrespective of packet size. In the following section, we explore the degree to which SoTA transcievers can be duty-cycled to achieve CaIN-like power numbers.

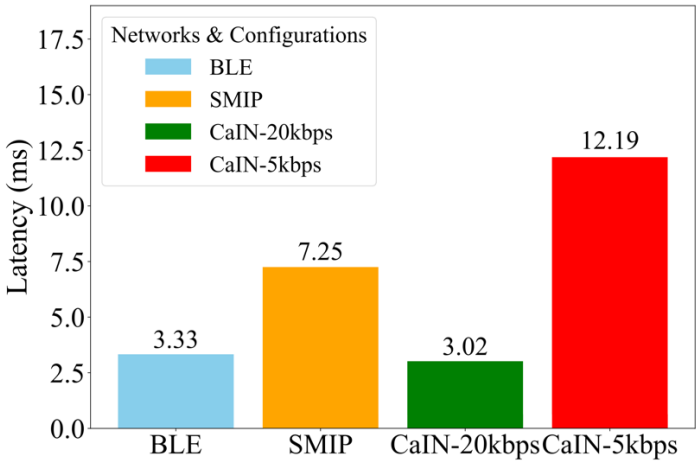

Figure 7.8: Latency for CaIN vs SOTA radios at 100% duty cycle.

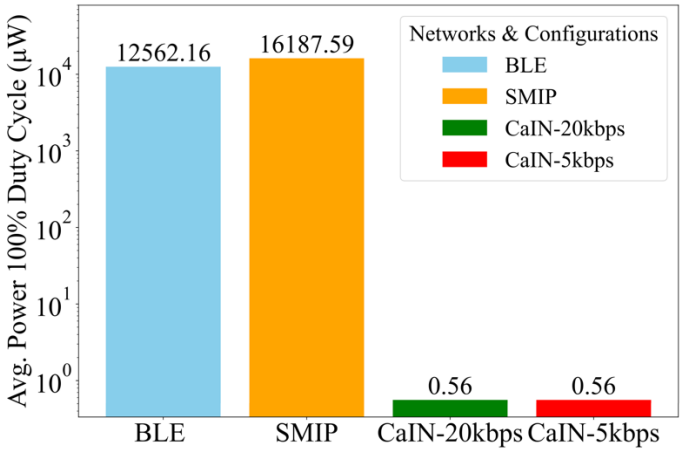

Figure 7.9: Average power consumption for SOTA radios and CaIN at 100% duty cycle.

There is a relationship between average receiver power consumption and duty cycle. As the different transceivers in our evaluation have different slot lengths, period (i.e. latency) is used in place of duty cycle. BLE power consumption converges with a CaIN receiver for a duty cycle period of around 100 seconds, while SMIP power consumption converges for a duty cycle period of around 200 seconds. This is orders of magnitude longer than CaIN's worst case latency of 12ms. We believe that this makes a strong case for CaIN in scenarios that demand low power operation while maintaining low latency for stochastic flows. Despite these promising early results, in its current form, CaIN has a number of limitations. We discuss these in the following section.

Limitations: We believe that CaIN offers a novel solution to supporting low latency flows on low power wireless networks. However, the approach has a number of limitations as follows:

1. *Range:* the most obvious and fundamental limitation of CaIN is the short range of near-field communication. While we believe that there is scope for incremental range extensions for CaIN, these will always be less than offered by standard far-field radios.

2. *Speed:* while the link-layer throughput of CaIN is competitive with conventional low power wireless networks at 4.9 to 19.8kbps, the low frequency of the carrier signal, which is necessary to achieve a reasonable near-field range, limits data rates in comparison to higher frequency transmitters.

3. *Configurability*: To an extent, the results presented in this deliverable present a false dichotomy between transmitter power, speed and range. At very close ranges, higher transmission speeds are possible and, as can be inferred from **Figure 7.5**, a variable power amplifier is likely to result in significant transmitter power savings at intermediate ranges.

4. *Sub-optimal front end*: The analog front end of CaIN lacks both impedance matching and band-pass filtering. This implies that significant performance gains may yet be achieved by addressing these shortcomings. The results presented

in this deliverable should therefore be viewed as close to a worst-case for the proposed technology.

5. *Limited testbed evaluation:* We are confident that the latency results presented will generalise to larger networks as the observed data were essentially deterministic, however a longitudinal study of a larger scale CaIN test-bed will add vital insights into reliability and robustness that are not captured in the current study.

7.5 Related Work

CaIN improves the energy-latency trade-offs that characterise low-power networks by introducing the paradigm of near-field mesh networking to support low latency, aperiodic traffic. As will be discussed, several contemporary research lines improve the same trade-off, often introducing some form of network-wide coordination that is highly contingent on existing, far-field PHYsical-layer (PHY) design. CaIN's performance advantages on top of an arguably more primitive network layer hence demonstrate the necessity of a new physical- layer foundation.

Efficiently Disseminating Data: In duty-cycled networks that route packets along multiple hops, nodes can conserve energy by only waking up when they are part of the routing path. Likewise, if multiple forwarding paths cover a given region of space, the energy burden of active listening can be divided between possible forwarders. Low-power mesh networks hence leave considerable space for cross-layer optimization [35]. Staffetta [36], for instance, adjusts radio duty cycles in a routing topology-aware fashion. Likewise, so-called Coalesced Intermittency [37] desynchronises the duty cycles of intermittently powered nodes that are deployed close together to provide the illusion of near-continuous availability. Reconciling such opportunistic medium access control and routing with ultra-low power radio design to further improve energy-latency trade-offs [38] is nevertheless nontrivial: simulation results reveal that principles of "good" radio design subject to microwatt-level power constraints may vary with the routing topology [39] due to the interaction between radio sensitivity and medium access control. Similarly, fast coordinated responses to environmental changes require nodes to have

synchronised rather than desynschronised duty cycles [40]. As the runtime optimisation of such conflicting design criteria and mutually interacting network parameters is a considerable challenge itself [41], this work leaves the study of such cross-layer complexities to future work, instead demonstrating that the physical-layer assumptions on which current studies are based may not be fully justified.

A recent line of research sidesteps the complex interactions between radio design and routing by replacing the latter with Synchronous Transmissions (ST), a technique that efficiently disseminates data through network-wide broadcasts on top of highly parallelised link-layer anycasts, exploiting physical-layer properties to avoid interference between concurrent transmissions [12]. Still, reconciling such efficient dissemination with aperiodic traffic and low average power consumption is nontrivial, since the choice of PHY determines power consumption, the degree of synchronisation enabled by broadcast wake-up signals, and the degree of synchronisation required for ST (cf. [42, 43, 17, 44]). Zippy [17] solves these interconnected problems by introducing an ST-based 434 MHz (far-field) network design based on On-Off Keying, delivering a 2-byte packet across 3 hops in roughly 40 milliseconds, with a per-node power consumption of 8 µW when receiving and 45 mW when transmitting. In scenarios where communication is sparse in time and receiver behavior dominates average power consumption, this work hence demonstrates order-of- magnitude power reductions while further reducing latency through appropriate PHY design. Similarly, Radunovic et al. have argued that the adoption of mesh networks relying on carrier frequencies in the 100s of MHz would enable new transmission modalities (i.e. full-duplex wireless) and simplifies cross-layer design concerns [45].

Sleeping Early and Often: Low-power, low-latency wireless networking requires that nodes sleep as often as possible. Nodes should hence avoid waiting for data that will not arrive and coordinate channel access to prevent wasted listening time due to contention and interference [46]. Blitz [47] hence coordinates access to an ST- based communication channel, starting from 434 MHz WuRs, thus mitigating contention. Similarly, Crystal [48] suppresses irrelevant transmissions; as sensor data

evolves according to predictable patterns, data collection applications can suffice with scarce, aperiodic, ST-based updates that indicate when sensor data deviates from its predicted evolution. While such higher-level protocol considerations are not addresses in this work recent results highlight the importance of appropriate PHY design in this context as swell, albeit without considering the impact of always-on radios: ST-like broadcasting of ultra-wideband preambles, combined with preamble sampling, energy-efficiently signal the presence or absence of data and hence reduce the energy consumption of the protocols mentioned in this paragraph by enabling early termination of their execution [49].

Complementary PHY Techniques: State-of-the-art low-power, low-latency mesh networks incorporate several additional physical-layer mechanisms that could further improve this CaINs approach. Blitz [47], for example, shows that on-off keyed wake-up signals exhibit temporal features that enable the fast detection of spurious wakeups using simple classifiers. Structured or modulated preambles can likewise improve performance through hardware-based false-positive wake-up mitigation [50] and thereby help overcome interference [17]. A recent stream of backscatter innovation also shows that CaIN-like envelope detection may achieve hundreds of metres of range under more challenging signal-to-noise regimes than tackled in this work by using modulated transmissions, but these efforts for now remain limited to transmitters and receivers that consume on the order a single watts and tens of microwatts, respectively [51, 52, 53]. Nevertheless, nearly passive radios should not be expected to achieve data rates and ranges comparable to those of conventional radios due to their lower sensitivity [39]: several of the aforementioned platforms rely on heterogeneous, wake-up radio designs, which switch to a second, faster, more power consuming radio mode after a wake-up phase [54, 56]. Such dual-radio schemes could be built on top of CaIN as future work, though CaIN demonstrates that near-field mesh networking already outperforms arguably more optimised, far-field approaches, while introducing considerably less hardware complexity.

For some applications, deployment specifics simplify low-power, low-latency mesh network design. The spatio- temporally correlated way in which seismic events

present themselves at distributed sensor nodes, for example, helps coordinate wake-up and data dissemination to further improve power-latency trade-off in mesh networks [54]. Likewise, deployment-specific constraints on network topology simplify the cross-layer design concerns described previously [57].

7.6 Future Work

Our future work will proceed along three fronts. A performance-oriented track will explore how to increase the efficiency of the CaIN transmitter, and thereby lower the power consumption of message transmission. While the current hardware implementation is only effective to 20 kbps, initial evidence from lab experiments suggests that link-level data rates of up to 100kbps may be possible through a redesign of the data slicer sub-system.

The protocol-oriented track of our future work will focus on data link and network layer protocols that exploit the unique near-free listening property of CaIN. The current single-controller bus architecture is rather restrictive and we intend in the future to investigate multi-controller approaches, which will allow any CaIN node to initiate communication. This will necessitate the extension of the CaIN data link layer with contention management techniques such as Carrier Sense Medium Access (CSMA). We will also investigate how existing approaches such as the Routing Protocol for Low-Power and Lossy Networks (RPL) [33] and Ad-Hoc On Demand Distance Vector (AODV) [34] can be applied for static and mobile scenarios respectively.

Finally, in an engineering-oriented track, we will first scale-up the CaIN testbed to 10s of nodes and then explore the extent to which CaIN may be brought into compliance with certification standards for the 40MHZ VHF radio control band. We also envisage the creation of a dedicated packet-based transceiver module, which offers a simplified interface to the application MCU and a greater degree of dynamic transceiver configuration.

8. Conclusion

This deliverable has summarized the outcomes of Task 2.2, which focuses on realizing a Radio Frequency (RF) version of the ZeroWire protocol from KU Leuven. That is to say, a low latency and yet low power wireless network that behaves, as much as possible, like a wired network. To date this has led to three scientific contributions, each with a matching publication (in the case of CaIN, the paper remains under submission). The technical contributions are summarized in **Table 8.1** below.

Table 8.1: Technical contributions

Technical Outcome	Type	TRL	Key Performance Data
mmWave ZeroWire Analog (Asynchronous Burst Link - ABL) [2]	Simulation (source code and data set)	TRL 2	99% reliability and 1ms latency over 8 hops. Data rate of 64 kbps.
ISM ZeroWire Analog [3]	Simulation (source code and data set)	TRL 2	99% reliability and latency of a few milliseconds over multiple hops. The latency of 0.3% increase per extra hop. Data rate of 100 kbps.
VHF ZeroWire Analog (Capacitive and Inductive Network - CaIN) [4].	Open hardware and software design.	TRL 3	95% reliability, sub-µW receiver power, 20kbps goodput and under 50ms latency for a 10-hop network.

In **Section 5**, we introduced ABL, an innovative protocol for IoT mesh networks. ABL is the first radio frequency approach to symbol synchronous communication and includes a novel differential pulse modulation scheme. We evaluate the latency and reliability

performance in a 9-node network with a maximum of 8 hops to transmit 32-bit packets. We conducted simulations at a sampling frequency of 25 MHz, revealing that the latency for all nodes could be reduced to less than 650µs. Our Differential Pulse (DP) modulation scheme ensures a system reliability of 99% even at a Signal-to-Noise Ratio (SNR) as low as -10 dB. The numerical results substantiate that our proposed ABL presents an effective strategy for achieving low latency and high reliability, while enhancing mm-wave communication coverage via a spatially redundant multi-hop mesh network architecture. This research thus provides a potential solution for the critical design challenge of mm-wave transceivers, fulfilling the stringent low latency and robustness requirements essential for critical industrial network applications.

In **Section 6**, we proposed an original symbol-synchronous radio design that aims at providing reliable low-latency communication in challenging WSN scenarios, such as industrial IoT, or robotic swarms, within the 2.4 GHz ISM band. In contrast to traditional WSN approaches, the proposed scheme utilizes concurrent transmissions and enables symbol-by-symbol relaying, thus providing wire-like delays through a wireless medium. Our early results show that the proposed design has extremely high potential. With a low transmission power of 0 dBm, it allows transmitting the signals with BER of 0.04% through in a 4-by-4 node network topology over a range of 250 m from the initiator to the furthest receiver (i.e., a grid distance of 60 m). In addition, in a 7-by-7 node network topology with a grid distance of 75 m, the current design still achieves to transmit data with BER of 5%. As for latency, it can achieve an end-to-end latency of about 5 milliseconds for a 512-bit packet and about 0.64 milliseconds for a 64-bit packet. Furthermore, while store-and-forward networks scale linearly with the number of hops in terms of latency, our results have shown significantly better scaling behavior, with a latency increase of only 0.3% per hop. It should be noted that there is a trade-off between the data rate and the experienced bit error rate. A higher data rate implies lower symbol duration, which may result in inter-symbol interference for networks with a large number of nodes. The optimal selection of symbol duration will be considered in our future work.

In **Section 7**, we introduced CaIN, a novel approach to low-power mesh networking which depends upon near field effects to dramatically reduce the power of idle listening (558nW). Using CaIN, we introduce a new concept, near field mesh networking, which offers a unique combination of low power and low latency networking. This is primarily due to CaIN'ss always-on operation, which maintains a per-hop latency under 12ms in all configurations; close to Bluetooth Low Energy (BLE) or 802.15.4/Zigbee at a 100% duty cycle, in which mode these protocols consume thousands of times more energy. The CaIN transmitter offers two power levels. At 23.5 mW transmit power, CaIN offers 4.9kbps of link layer throughput, 12.2ms per-hop latency and 7m range with over 80% reliability. At 201.7mW transmit power, CaIN offers 19.8kbps of link layer throughput, 3.02 ms per-hop latency and 17m of range at 80% reliability. With the exception of range, these characteristics compare favorably to state-of-the-art networks. We believe that, despite CaINs simple and low-cost transceiver design, the achieved performance envelope points towards a bright future for near-field mesh networking. Stepping back from technical specifics, we believe that there CaIN offers one more piece of evidence that novel, clean slate radio designs can significantly improve the performance of resource-constrained wireless networking.

Next Steps:

- **ABL:** Moving from simulation to real-world implementation using either commodity mmWave radars or software defined radios. We are currently experimenting with both approaches.
- **ISM Zero-Wire:** Implementing the simulated network into hardware platform like software defined radio. Considering the interference mitigation techniques to optimize the system robustness.
- **CaIN:** Accepted at EWSN. Reworking the CaIN reference hardware to match the common project form factor and integrating the board with one or more of the swarm robotics platforms in use by the consortium. Exploring deployment of CaIN in the project demonstrators.

List of Publications

- Bingwu Fang, Jonathan Oostvogels, Xinlei Liu, Andrey Belogaev, Sam Michiels, Jeroen Famaey, Danny Hughes, **ABL: Leveraging Millimeter Wave Pulses for Low Latency IoT Networking**, CPS-IoT Week '24: Cyber-Physical Systems and Internet of Things Week 2024, CPS-IoT Week '24: Proceedings of Cyber-Physical Systems and Internet of Things Week 2024, Hong Kong, China, May 13-16, 2024
- Xinlei Liu, Andrey Belogaev, Jonathan Oostvogels, Bingwu Fang, Danny Hughes, Maarten Weyn, and Jeroen Famaey, **Low-latency Symbol-Synchronous Communication for Multi-hop Sensor Networks**, EuCNC&6G Summit, Antwerp, Belgium, June 3-6, 2024
- Mengyao Liu, Bingwu Fang, Jonathan Oostvogels, Sam Michiels, Xinlei Liu, Andrey Belogaev, Jeroen Famaey and Danny Hughes, **CaIN: Low Power and Low Latency VHF Mesh Networking**, accepted for publication at the 21st International Conference on Embedded Wireless Systems and Networks (EWSN '24).

Bibliography

[1] Jonathan Oostvogels, Fan Yang, Sam Michiels, Danny Hughes, **Zero-Wire: A Deterministic and Low-Latency Wireless Bus Through Symbol-Synchronous Transmission of Optical Signals**, (eds. Jin Nakazawa, Polly Huang), SenSys '20: Proceedings of the 18th Conference on Embedded Networked Sensor Systems, SenSys '20: Proceedings of the 18th Conference on Embedded Networked Sensor Systems, pages 164-178, Virtual Event, Japan, November 16-19, 2020

[2] Bingwu Fang, Jonathan Oostvogels, Xinlei Liu, Andrey Belogaev, Sam Michiels, Jeroen Famaey, Danny Hughes, **ABL: Leveraging Millimeter Wave Pulses for Low**

Latency IoT Networking, CPS-IoT Week '24: Cyber-Physical Systems and Internet of Things Week 2024, CPS-IoT Week '24: Proceedings of Cyber-Physical Systems and Internet of Things Week 2024, Hong Kong, China, May 13-16, 2024

[3] Xinlei Liu, Andrey Belogaev, Jonathan Oostvogels, Bingwu Fang, Danny Hughes, Maarten Weyn, and Jeroen Famaey, **Low-latency Symbol-Synchronous Communication for Multi-hop Sensor Networks**, EuCNC&6G Summit, Antwerp, Belgium, June 3-6, 2024

[4] Mengyao Liu, Bingwu Fang, Jonathan Oostvogels, Sam Michiels, Xinlei Liu, Andrey Belogaev, Jeroen Famaey and Danny Hughes, **CaIN: Low Power and Low Latency VHF Mesh Networking**, accepted for publication at the 21st International Conference on Embedded Wireless Systems and Networks (EWSN '24).

[5] Devan, P. A. M., Hussin, F. A., Ibrahim, R., Bingi, K., and Khanday, F. A. **A survey on the application of wirelesshart for industrial process monitoring and control**. Sensors 21, 15 (2021).

[6] Vilajosana, X., Watteyne, T., Chang, T., Vucinic, M., Duquennoy, S., and Thubert, P. **IETF 6TISCH: A tutorial**. IEEE Communications Surveys & Tutorials 22, 1 (2020), 595–615.

[7] Watteyne, T., Doherty, L., Simon, J., and Pister, K. **Technical overview of smartmesh IP**. In 2013 Seventh International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing (2013), pp. 547–551.

[8] R.Rondo ́n,A.Mahmood,S.Grimaldi,andM.Gidlund. **Understanding the Performance of Bluetooth Mesh: Reliability, Delay, and Scalability Analysis.** IEEE Internet of things journal, vol. 7, no. 3, pp. 2089–2101, 2019.

[9] "Industrial Networks–Wireless Communication Network and Communication Profiles–WIA-PA," International Electrotechnical Commission, IEC Standard PAS 62601, 2015, available: https://webstore.iec.ch/publication/23902.

[10] K. Montgomery, R. Candell, Y. Liu, and M. Hany. **Wireless user requirements for the factory workcell.** Department of Commerce, National Institute of Standards and Technology, 2020.

[11] M. Baddeley, C. A. Boano, A. Escobar-Molero, Y. Liu, X. Ma, V. Marot, U. Raza, K. Ro ̈mer, M. Schuß, and A. Stanoev. **Understanding Concurrent Transmissions: The Impact of Carrier Frequency Offset and RF Interference on Physical Layer Performance.** ACM Transactions on Sensor Networks, 2023.

[12] M. Zimmerling, L. Mottola, and S. Santini. **Synchronous transmissions in low-power wireless: A survey of communication protocols and network services.** ACM Computing Surveys (CSUR), vol. 53, no. 6, pp. 1–39, 2020.

[13] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh. **Efficient network flooding and time synchronization with glossy.** In Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks. IEEE, 2011, pp. 73–84.

[14] D. Lobba, M. Trobinger, D. Vecchia, T. Istomin, G. P. Picco et al. **Concurrent Transmissions for Multi-hop Communication on Ultra-wideband Radios.** in EWSN, 2020, pp. 132–143.

[15] B. A. Nahas, A. Escobar-Molero, J. Klaue, S. Duquennoy, and O. Landsiedel. **BlueFlood: Concurrent transmissions for multi-hop Bluetooth 5 — Modeling and evaluation.** ACM Transactions on Internet of Things, vol. 2, no. 4, pp. 1–30, 2021.

[16] X. Xia, N. Hou, Y. Zheng, and T. Gu. **PCUBE: scaling LORA concurrent transmissions with reception diversities.** ACM Transactions on Sensor Networks, vol. 18, no. 4, pp. 1–25, 2023.

[17] F. Sutton, B. Buchli, J. Beutel, and L. Thiele. **Zippy: On-demand Network Flooding.** In Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems, 2015, pp. 45–58.

[18] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez. **Millimeter wave mobile communications for 5G cellular: It will work!** IEEE access, vol. 1, pp. 335–349, 2013.

[19] K. Yamamoto, F. Ichihara, K. Hasegawa, M. Tukuda, and I. Omura. **60 GHz wireless signal transmitting gate driver for IGBT.** In 2015 IEEE 27th International Symposium on Power Semiconductor Devices & IC's (ISPSD). IEEE, 2015, pp. 133–136.

[20] T. Watteyne. **Crystal-free architectures for smart dust and the industrial IoT.** In 2020 7th International Conference on Internet of Things: Systems, Management and Security. IEEE, 2020, pp. 1–1.

[21] J. M. Jornet and I. F. Akyildiz. **Femtosecond-long pulse-based modulation for terahertz band communication in nanonetworks.** IEEE Transactions on Communications, vol. 62, no. 5, pp. 1742–1754, 2014.

[22] Mi Wave, **"840EF-17/387, E-band, WR-12, 17 dBm Output Power, 60 GHz-90 GHz Wide Band Voltage Controlled Sources (VCO),"** [On-line]. Available: https://www.miwv.com/voltage-controlled-oscillators-wide-band-sources-vco-e-band-60ghz-to-90ghz/, 2007.

[23] Texas Instruments, **"IWR6843 Single-Chip 60-GHz to 64-GHz mmWave Sensor Data Sheet,"** 2023, [Online]. Available: http://www.ti.com/lit/ds/symlink/iwr6843.pdf.

[24] Analog Devices, **"HMC1144-DIE 35 GHz to 70 GHz, GaAs, pHEMT, MMIC, Medium Power Amplifier,"** [Online]. Available: https://www.analog.com/media/en/technical-documentation/datasheets/hmc1144.pdf, 2015.

[25] A. W. Mbugua, Y. Chen, L. Raschkowski, L. Thiele, S. Jaeckel, and W. Fan. **"Review on Ray Tracing Channel Simulation Accuracy in Sub-6 GHz Outdoor Deployment Scenarios."** IEEE Open Journal of Antennas and Propagation, vol. 2, pp. 22–37, 2021.

[26] Buettner, M., Yee, G. V., Anderson, E., and Han, R. **X- mac: a short preamble mac protocol for duty-cycled wireless sensor networks**. In Proceedings of the 4th International Conference on Embedded Networked Sensor Systems (New York, NY, USA, 2006), SenSys '06, Association for Computing Machinery, p. 307–320.

[27] Polastre, J., Hill, J., and Culler, D. **Versatile low power media access for wireless sensor networks.** In Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems (New York, NY, USA, 2004), SenSys '04, Association for Computing Machinery, p. 95–107.

[28] Martinussen, E. S., and Arnall, T. **Designing with RFID**. In Proceedings of the 3rd International Conference on Tangible and Embedded Interaction (New York, NY, USA, 2009), TEI '09, Association for Computing Machinery, p. 343–350.

[29] Piyare, R., Murphy, A. L., Kiraly, C., Tosato, P., and Brunelli, D. **Ultra low power wake-up radios: A hardware and networking survey**. IEEE Communications Surveys & Tutorials 19, 4 (2017), 2117–2157.

[30] Bdiri, S., and Derbel, F. **A nanowatt wake-up receiver for industrial production line.** In 2014 IEEE 11th International Multi-Conference on Systems, Signals & Devices (SSD14) (2014), pp. 1–6.

[31] Marinkovic, S., and Popovici, E. **Ultra low power signal oriented approach for wireless health monitoring**. Sensors 12, 6 (2012), 7917– 7937.

[32] Magno, M., Jelicic, V., Srbinovski, B., Bilas, V., Popovici, E., and Beninin, L. **Design, implementation, and performance evaluation of a flexible low-latency nanowatt wake-up radio receiver**. IEEE Transactions on Industrial Informatics 12, 2 (2016), 633–644.

[33] Winter, T. **RFC 6550: RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks** — datatracker.ietf.org. https://datatracker. ietf.org/doc/html/rfc6550. [Accessed 02-05-2024].

[34] Perkins, C., and Royer, E. **Ad-hoc on-demand distance vector routing**. In Proceedings WMCSA'99. Second IEEE Workshop on Mobile Computing Systems and Applications (1999), pp. 90–100.

[35] Ghribi, M., and Meddeb, A. **Survey and taxonomy of mac, routing and cross layer protocols using wake-up radio**. Journal of Network and Computer Applications 149 (2020), 102465.

[36] Cattani, M., Loukas, A., Zimmerling, M., Zuniga, M., and Langendoen, K. **Staffetta: Smart duty-cycling for opportunistic data collection**. In Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems (2016), pp. 56–69.

[37] Majid, A. Y., Schilder, P., and Langedoen, K. **Continuous sensing on intermittent power**. In 2020 19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN) (2020), pp. 181–192.

[38] Oller, J., Demirkol, I., Casademont, J., Paradells, J., Gamm, G. U., and Reindl, L. **Has time come to switch from duty- cycled mac protocols to wake-up radio for wireless sensor networks?** IEEE/ACM Transactions on Networking 24, 2 (2015), 674–687.

[39] Longman, E., El-Hajjar, M., and Merrett, G. V. **Multihop networking for intermittent devices**. In Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems (New York, NY, USA, 2023), SenSys '22, Association for Computing Machinery, p. 878–884.

[40] Geissdoerfer, K., and Zimmerling, M. **Bootstrapping battery-free wireless networks: Efficient neighbor discovery and synchronization in the face of intermittency.** In 18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21) (2021), pp. 439–455.

[41] Zimmerling, M., Ferrari, F., Mottola, L., Voigt, T., and Thiele, L. **pTunes: runtime parameter adaptation for low-power MAC protocols**. In Proceedings of the 11th International Conference on Information Processing in Sensor Networks (New York, NY, USA, 2012), IPSN '12, Association for Computing Machinery, p. 173–184.

[42] Baddeley, M., Boano, C. A., Escobar-Morelo, A., Liu, Y., Ma, X., Raza, U., Romer, K., Schuss, M., and Stanoev, A. **The impact of the physical layer on the performance of concurrent transmissions**. In 2020 IEEE 28th International Conference on Network Protocols (ICNP) (2020), pp. 1–12.

[43] Liu, M., Oostvogels, J., Michiels, S., Joosen, W., and Hughes, D. **BoboLink: low latency and low power communication for intelligent environments**. In 2022 18th International Conference on Intelligent Environments (IE) (2022), IEEE, pp. 1–4.

[44] Yang, F., Oostvogels, J., Michiels, S., AND Hughes, D. **Achieving deterministic and low-latency wireless connection with Zero- Wire: demo abstract**. In Proceedings of

the 18th Conference on Embedded Networked Sensor Systems (SenSys) (2020), ACM, pp. 591–592.

[45] Radunovic, B., Gunawardena, D., Key, P., Proutiere, A., Singh, N., Balan, V., and Dejean, G. **Rethinking indoor wireless mesh design: Low power, low frequency, full-duplex**. In 2010 Fifth IEEE Workshop on Wireless Mesh Networks (2010), pp. 1–6.

[46] Sutton, F., Da Forno, R., Gschwend, D., Gsell, T., Lim, R., BeutelL, J., and Thiele, L. **The design of a responsive and energy- efficient event-triggered wireless sensing system**. In EWSN (2017), pp. 144–155.

[47] Sutton, F., Forno, R. D., Beutel, J., and Thiele, L. **Blitz: Low latency and energy-efficient communication for event-triggered wireless sensing systems**. ACM Trans. Sen. Netw. 15, 2 (mar 2019).

[48] Istomin, T., Murphy, A. L., Picco, G. P., and Raza, U. **Data prediction + synchronous transmissions = ultra-low power wireless sensor networks**. In Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM (New York, NY, USA, 2016), SenSys '16, Association for Computing Machinery, p. 83–95.

[49] Soprana, E., Trobinger, M., Vecchia, D., and Picco, G. P. **Network on or off? instant global binary decisions over uwb with flick**. In Proceedings of the 22nd International Conference on Information Processing in Sensor Networks (New York, NY, USA, 2023), IPSN '23, Association for Computing Machinery, p. 261–273.

[50] Petrioli, C., Spenza, D., Tommasino, P., and Trifiletti, A. **A novel wake-up receiver with addressing capability for wireless sensor nodes**. In 2014 IEEE International Conference on Distributed Computing in Sensor Systems (2014), pp. 18–25.

[51] Guo, X., Shangguan, L., He, Y., Jing, N., Zhang, J., Jiang, H., and Liu, Y. Saiyan: **Design and implementation of a low-power demodulator for {LoRa} backscatter systems**. In 19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22) (2022), pp. 437–451.

[52] Li, S., Zheng, H., Zhang, C., Song, Y., Yang, S., Chen, M., Lu, L., and Li, M. **Passive DSSS: Empowering the downlink communication for backscatter systems**. In 19th

USENIX symposium on networked systems design and implementation (NSDI 22) (2022), pp. 913–928.

[53] Song, Y., Lu, L., Wang, J., Zhang, C., Zheng, H., Yang, S., Han, J., and Li, J. **µMote: enabling passive chirp de-spreading and µW-level Long-Range downlink for backscatter devices**. In 20th USENIX symposium on networked systems design and implementation (NSDI 23) (2023), pp. 1751–1766.

[54] Biri, A., Forno, R. D., Gsell, T., Gatschet, T., Beutel, J., and Thiele, L. **STeC: Exploiting spatial and temporal correlation for event-based communication in wsns**. In Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems (New York, NY, USA, 2021), SenSys '21, Association for Computing Machinery, p. 274–287.

[55] Thangarajan, A. S., Nguyen, T. D., Liu, M., Michiels, S., Yang, F., Man, K. L., Ma, J., Joosen, W., and Hughes, D. **Static: Low frequency energy harvesting and power transfer for the internet of things**. Frontiers in Signal Processing 1 (2022).

[56] Piyare, R., Murphy, A. L., Magno, M., and Benini, L. **On-Demand LoRa: Asynchronous TDMA for Energy Efficient and Low Latency Communication in IoT**. Sensors 18, 11 (2018).

[57] Karvonen, H., Petajajarvi, J., Iinatti, J., Hamalainen, M., and Pomalaza-Raez, C. **A generic wake-up radio based mac protocol for energy efficient short range communication**. In 2014 IEEE 25th Annual International Symposium on Personal, Indoor, and Mobile Radio Communication (PIMRC) (2014), pp. 2173–2177.

List of Tables

List of Figures

Figure 7.6	Network densification affects signal strength due to near-field coupling, though impacts are manageable
Figure 7.7	End-to-end Latency for networks of 1 to 3 hops
Figure 7.8	Latency for CaIN vs SOTA radios at 100% duty cycle
Figure 7.9	Average power consumption for SOTA radios and CaIN at 100% duty cycle