



HAL
open science

A Time Synchronized Multi-Hop Mesh Network with Crystal-Free Nodes

Filip Maksimovic, Austin Patel, David Burnett, Thomas Watteyne, Kristofer S J Pister

► **To cite this version:**

Filip Maksimovic, Austin Patel, David Burnett, Thomas Watteyne, Kristofer S J Pister. A Time Synchronized Multi-Hop Mesh Network with Crystal-Free Nodes. GLOBECOM 2023 - IEEE Global Communications Conference, IEEE, Dec 2023, Kuala Lumpur, Malaysia. hal-04394968

HAL Id: hal-04394968

<https://inria.hal.science/hal-04394968>

Submitted on 15 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

A Time Synchronized Multi-Hop Mesh Network with Crystal-Free Nodes

Filip Maksimovic¹, Austin Patel², David Burnett³, Thomas Watteyne¹, Kristofer S.J. Pister²

¹ Inria, Paris, France {filip.maksimovic, thomas.watteyne}@inria.fr

² Berkeley Sensor and Actuator Center, University of California, Berkeley, USA {austinpate1, ksjp}@berkeley.edu

³ Portland State University, Portland, OR, USA dburnett@ece.pdx.edu

Abstract—In this work we propose and demonstrate a protocol for a time synchronized channel hopping mesh network for wireless transceivers that use exclusively imprecise and inaccurate on-chip oscillators. This protocol is built on an IEEE 802.15.4 physical layer radio that enables interoperability with protocols such as 6TiSCH or Thread. A calibration-bootstrapped multi-hop mesh network is demonstrated with a single crystal-enabled node acting as the root.

The protocol is designed to create a multi-hop mesh while compensating noisy and drifting oscillators and timers. With a 4 s synchronization period, an experimental implementation of the network maintains, in the worst case, 1.8 ms 3σ absolute time synchronization and 820 μ s 3σ hop-to-hop synchronization across four hops, under ambient environmental conditions. The resistance to environmental variation is tested by varying one node’s supply voltage. With time and frequency feedback from received packets, the node maintains this synchronization with a supply variation of 2.5 mV/s, which is equivalent to a temperature variation of 10°C/min with a packet rate of 0.5 Hz.

Index Terms—Internet of Things, Low-Power Wireless Mesh Networks, Smart Dust

I. INTRODUCTION

A crystal reference in a wireless node typically serves two purposes. First, as a reference in a phase locked loop (PLL) that synthesizes the radio frequency oscillation that dictates channel. Second, as a timer (often 32 kHz) used as a reference for the processor clock and as a system timer for networking or as a watchdog timer. Although this is an active area of research, most commercial wireless nodes have two crystals, one in the MHz range used for the RF PLL, one at 32 kHz that serves as the low-frequency timer.

Crystals are an accurate, precise, and environmentally-invariant frequency reference. Even the lowest performance crystals offer better than 40 parts per million (ppm) of error. This 40 ppm figure appears in both the IEEE 802.15.4 [1] and Bluetooth [2] specifications for time and frequency specificity. If a transceiver does not have a crystal, the node has little understanding of what absolute time and frequency are, and has noisy oscillators that can change in the presence of environmental variation. In this work, we propose and measure a calibrated multi-hop mesh network with a single crystal-enabled transceiver acting as a root node. We demonstrate that crystal references are not necessary for low power mesh networks, enabling further miniaturization and power reduction of IoT devices.

The first reason for removing external crystals is to reduce power consumption. A 32 kHz crystal consumes little current, typically less than 1 μ A [3]. The RF crystal and corresponding PLL, on the other hand, can burn significant power. Even in recent work, such as [4], the power consumption of the frequency synthesis alone is over 1 mW from a 0.5 V supply, which is comparable to the power consumption of the entire system-on-chip used in this paper. The high power consumption is because of high-frequency divider and high-speed phase comparison, both of which can consume significant energy, albeit over very short periods of time if the radio is heavily duty-cycled. The second reason is physical volume. If the crystal is removed, the entirety of a wireless transceiver can be integrated onto a single CMOS die that requires only a power source and an antenna to communicate. Such devices already exist, most notably RFID tags, but these devices are not capable of networking because they can only communicate with a dedicated reader in an asymmetrical link. This makes them unsuitable for independent communication necessary for mesh networking. The operation of an independent transceiver with no external components and only three external connections (power, ground, and antenna) brings us one step closer to a complete “Smart Dust” device: a single-chip wireless mote that can be extremely small and inexpensive.

To demonstrate the feasibility of crystal-free nodes in an industrial IoT style mesh network, this paper presents three contributions related to the calibration and compensation of imprecise frequency and time. First, a rapid low-energy calibration scheme that is used to calibrate each individual node so that it uses standards-specified channels in the 2.4 GHz ISM band. Second, a demonstration of a multi-hop mesh network that shows the impact of imprecise oscillators on guard time requirements and time synchronization. This network is briefly compared to existing crystal-enabled time synchronization protocols. Third, the use of periodic packets to maintain connectivity in the presence of rapid environmental variation. The device used for the experiments, the Single-Chip Micro-Mote (SC μ M) [5], in this work is shown in Fig. 1.

II. HARDWARE SPECIFICS

The pertinent specifications for time synchronization and channel hopping are as follows. First, a fully functional 802.15.4e transceiver with -83 dBm channel tuning capability

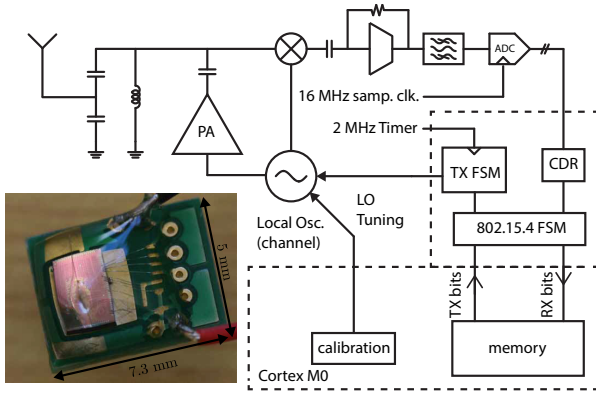


Fig. 1. Schematic and photograph of the single-chip micro-mote (SC μ M), the experimental platform used in this paper. The schematic shows the relevant on-chip oscillators: the channel-select oscillator, the transmit chipping clock, and the receive sampling clock.

of 90 kHz resolution and a range that covers the 2.4 GHz ISM band. Second, three low frequency timers: one at 2 MHz that sets the chipping rate of the transmitter, one at 16 MHz that sets the sampling clock of the receiver, and one at 500 kHz which is derived from the on-chip ARM Cortex-M0 and that is used for timing purposes. All of these oscillators are free-running: without an initial calibration and constant updates to that calibration, they will drift. More details about the hardware are available in a previous publication [5].

III. THE PROPOSED PROTOCOL

To evaluate the use of SC μ M for industrial IoT protocols such as 6TiSCH, we propose μ TSCH, a minimalistic implementation of Time Synchronized Channel Hopping (TSCH). TSCH is the core technique used in standards such as 6TiSCH and WirelessHART that are widely deployed in industrial IoT applications. The proposed protocol implements both Medium Access Control (MAC) and networking layers. Higher-level protocol responsibilities, like routing and scheduling, are pre-computed beforehand and hard coded. At the MAC layer, the protocol cuts time into 10 ms time slots, in which a node either transmits, receives, or sleeps, depending on its state and the schedule. The purpose of the protocol is to mimic the basic framework of existing wireless protocols that use the 802.15.4 physical layer while generating a consistent network clock whose time error will be used to determine the ability of crystal-free nodes to synchronize. The goal is to minimize energy consumption while maintaining connectivity in the network. A simplified diagram of the proposed protocol is shown in Fig. 2.

A root node, which in the experiments in this paper has a crystal, sends periodic enhanced beacons (EBs) that indicate the beginnings of a slot frame. These beacons include the network EB rate and a join window that a child can use to attempt to join in the same way as in [6].

In order for a leaf node to synchronize and join, it leaves its receiver on until it hears at least two consecutive EBs, and

uses that information to calibrate its 500 kHz RFTimer. In the typical operation of the protocol the node will then send a join request to its parent. In the acknowledgement of the join request, it will receive a dedicated uplink/downlink slot and a slot to transmit its own EBs so that further leaf nodes can join.

One novel aspect of this protocol is that any non-root node's schedule is relative to the timing of its parent's EB, so there is no notion of a "global" slot frame or universal time. The effect of this design decision is to minimize the amount of hardware a leaf node must have to support this protocol – the only thing that a leaf node needs to do is reset a timer upon receiving an EB from its parent. In addition, time and frequency information only passes from parent to child. If a node used one of its child nodes for timing information, it no longer necessarily guarantees synchronization to its parent because the nodes further from the root node could have desynchronized since the last update.

There are two significant differences between the work presented here and previously proposed protocols. The first is that this work primarily concerns relative, node-to-node time error, because this sets the minimum guard time necessary to keep the receiver on as little as possible to conserve energy. Similar previously presented protocols, such as [7] often consider absolute global time error throughout the network. It also uses both downlink and uplink timer information to attain this goal. And, while global time error is measured, we prioritize guard time minimization to minimize power consumption for node-to-node communication reliability. The second difference is that, as stated in [8], time error in crystal-enabled networks is dominated by frequency differences between the crystal oscillators in various nodes caused either by manufacturing variation or changes in environment, like temperature. In this work, with exclusively on-chip oscillators, time error from environmental (temperature and supply) variations under ambient conditions is lower than the accumulated time error (jitter) caused by the electronic noise in the oscillators.

To explore results from a worst-case network stability perspective, time synchronization occurs if and only if a beacon is received with successful CRC check. Although start frames without CRC check are not used for time synchronization, they are still used to determine whether a node is desynchronized, as will be described later. The challenge in incorporating more advanced time synchronization schemes is that the time error caused by free-running on-chip oscillators is orders of magnitude larger than the synchronization error in crystal-enabled networks.

In the experiments in this paper, a single channel is used so that time information can be recorded by a single receiver. However, the protocol does support channel hopping, and each chip's local oscillator settings for multiple channels are found during the calibration procedure. With support for a join window and a dedicated data slot for a node's communication to both parent and a single child, the protocol uses approximately 17.5 kB of instruction RAM on the chip's Cortex-M0 as compiled by Arm's Keil IDE.




Fig. 2. Timing diagram of the proposed protocol for the root and first two hops of the network. The enhanced beacons (EBs) are essentially equivalent to unacknowledged keepalive packets.

A. Jitter in On-Chip Oscillators

As previously mentioned, the goal of the protocol to maintain synchronization between parent and child with the minimum possible energy consumption. In order to do this, the “guard time,” or the time between turning on a child’s receiver before it expects a packet, should be minimized.

Assuming the source of noise is stationary, an oscillator’s jitter statistics will be a random walk. If the statistics of the noise are white, then the N^{th} cycle of oscillation will have statistics:

$$\sigma_N = \sqrt{N}\sigma_c \quad (1)$$

Where σ_c is the jitter of a single cycle. As described in [9] and in [10] after a certain integration period, the noise will transition from flat power spectral density to a $1/f$ power spectral density. In electronics this is referred to as flicker noise. If an oscillator’s noise is in the flicker regime, the variance of its random walk will trend as:

$$\sigma_N = N\sigma_c \quad (2)$$

The flicker “corner,” or the frequency at which the noise transitions from white to $1/f$, varies significantly from process to process. In modern CMOS processes, it can easily be in the MHz range, which means that any fully-on-chip CMOS oscillator will be dominated by flicker noise almost immediately. This is effectively confirmed by the measurements of the crystal-free node’s RFTimer, which is the oscillator used to synchronize time, as shown in Fig. 3.

Note that this is not a measurement of Allan variance. Rather, it is simply a measurement of time error as a function of the timer’s divide ratio. It is worth mentioning that if this oscillator were dominated by white noise, this curve would express a square-root relationship.

B. Calibration Bootstrapping

The very first step is to perform a calibration procedure so that a crystal-free node can find the correct channel settings to communicate with a crystal-enabled node, like the root node in the networks in this work.

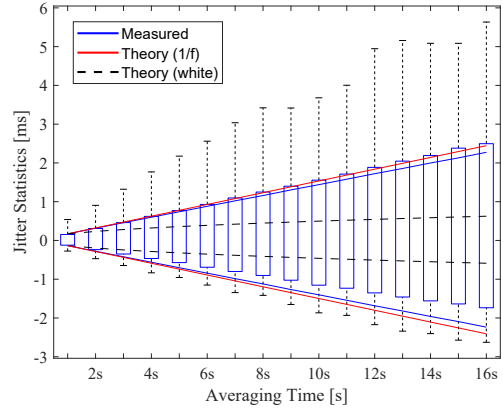


Fig. 3. Open-loop timer jitter showing the prevalent of $1/f$ “flicker” noise in the cycle-to-cycle time error of the oscillator.

A quick-calibration procedure for this particular chip was demonstrated in [11] but requires a bank of sixteen crystal-enabled transceivers (one for each 802.15.4 channel). We present a method that requires only one crystal-enabled transceiver. As implemented, the calibration procedure can only reliably be performed on one crystal-free radio at a time. The procedure is performed on each chip individually before the network is formed.

A time-domain illustration of the calibration is shown in Fig. 4. First, the crystal-enabled transceiver will continuously wait for packets on the lowest channel that needs to be found. The crystal-free mote blindly transmits packets on each of its DAC settings. The contents of the packets are the settings. If the crystal-enabled mote receives a packet, it stores the contents of the packet, for later processing, and starts a timeout timer. Once the timeout timer goes off, the receiver moves onto the next channel. This procedure is repeated until the final channel is reached and calibrated.

At this point, the crystal-enabled transceiver switches roles and begins transmitting beacons every 25 ms on the lowest channel. In between each beacon, it turns on its receiver on that same channel. The contents of its beacon packets are the transmit packets that it had previously received for this particular channel. When the crystal-free mote reaches the end of its DAC range, it also switches modes. It activates its receiver and sweeps its receive settings, waiting 30 ms at each one. If it receives a packet from the crystal-enabled transceiver, it responds using the transmit settings that it read from the received beacon. It continues sweeping its receive channel, receiving packets as it goes, until it determines that it is at the upper end of its receive bandwidth by observing the average downconverted intermediate frequency. When the crystal-free node is “finished” and has found appropriate receive settings for this channel, it sends an additional byte in the response packet that tells the crystal enabled node to move onto the next channel. Once again, this procedure is repeated until the crystal-free node has found its receive settings for every channel. After the crystal-free mote has swept through all of its




Fig. 4. Calibration of the crystal-free mote transmitter settings (top) and receive settings (bottom) across multiple channels with a single crystal-enabled radio (TI CC 2538).

receive settings, it transmits all of its codes, both transmit and receive, on the lowest channel. If a crystal-free node does not have a physical connection, it is still possible for a programmer to access the optimal settings for a given crystal-free mote at a given temperature and voltage. This procedure takes approximately 5 s per transmit channel and 20 s per receive channel, and the estimated energy consumption is 50 mJ per channel.

The frequency tuning characteristic in Fig. 4 exhibits a sawtooth waveform because there are overlapping codes in the capacitive digital-to-analog converter (DAC). In the calibration implemented in this work, multiple valid settings are found for both transmit and receive codes. Later on, when environmental (supply) variation is considered, this can extend the range of operation without rolling over DAC codes.

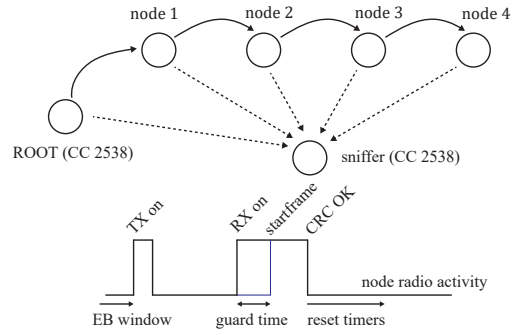


Fig. 5. The experimental setup to determine the quality of network time synchronization. For hop-to-hop or relative time error, statistics of the guard time (time difference between receiver on and received start frame) and statistics of absolute time are analyzed.

IV. EXPERIMENTAL NETWORK PERFORMANCE

In these experiments, we attempt to characterize network performance under the worst-case conditions. This occurs under two conditions: (a) the most hops, so that time error accumulates further from the root node, and (b) each node’s EB transmit-to-child window is set to occur in the slot just before (10 ms before) receiving the next EB from its parent. The second condition ensures near-maximum accumulation of time error in the EB transmission. To ensure our results could be extrapolated to an N-hop network, we created a four-hop network with one crystal-enabled node and four crystal-free nodes. The full experimental setup is shown in Fig. 5.

A. Ambient Conditions with a Crystal-Enabled Root Node

The protocol from Fig. 2 is implemented on four chips and each is given a unique MAC address. Received packets are screened by MAC address to enforce a multi-hop network. A logic analyzer is connected to two GPIO pins on each chip to analyze the time accuracy of an individual node’s transmission and reception of enhanced beacons. As previously mentioned, each node’s enhanced beacon transmission is set to occur 10 ms before it expects to receive an enhanced beacon from its parent to ensure that the maximum amount of time error is accumulated after synchronizing to its parent to mimic a “worst-case” scenario for synchronization. A plot of the statistics of the enhanced beacon transmission times, as received by the crystal-enabled sniffer (absolute time), and the time between a node’s RX on and startframe (guard time) are shown in Fig. 6.

The dotted lines in Fig. 6 represent the expected timer statistics at each node in the network. The variance of the transmitted packet, which we consider “absolute” time error, increases linearly with hop because it is simply a cumulative sum of time error of N (where N is the number of hops) uncorrelated random variables with $1/f$ power spectral density. On the other hand, the variance of the guard time, which is a “relative” time error between hops, appears to stay constant between nodes (besides the first node, which receives synchronization packets from a practically ideal crystal-

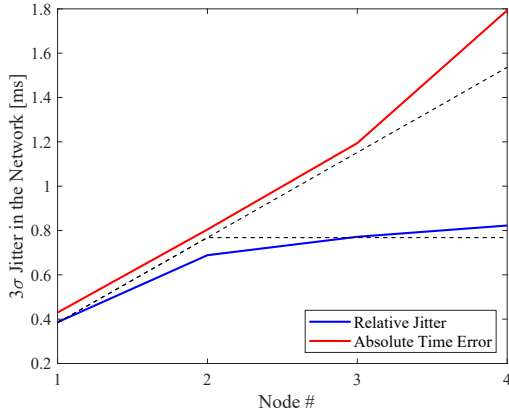


Fig. 6. Standard deviation of relative and absolute time error. Relative time error is between two adjacent hops and absolute time error is compared to the crystal reference (node 0).

enabled clock). This is expected because the node-to-node synchronization is simply the difference of two uncorrelated random variables. The constant relative jitter permits a fixed guard time throughout the network, regardless of the depth of the network. This is relevant for the implementation of time-synchronized mesh network stacks, such as WirelessHART, which rely on constant guard times. The apparent increase in variance vs. hop is caused by desynchronization events at the fourth node because of a low-SNR link. The absolute time error results as performed in the protocol from this work match the uncompensated results from [12] and [13], which indicates that if absolute time error were important to our network, there is room for improvement. The flattening out of the relative time error confirms the results in [14]. Even though the network has a single crystal, the time synchronization trends still follow the same results as crystal-enabled networks.

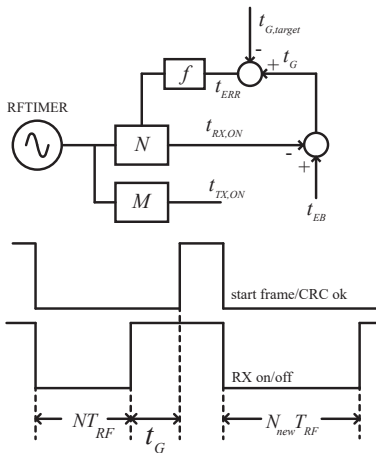


Fig. 7. Illustration of the feedback mechanism used to compensate the on-chip timer. The time axis of the radio activity is not to scale and is exaggerated to show the various times in the compensation strategy.

B. Over-the-air Supply Compensation

The previous experiments were performed under ambient conditions with ideal, regulated power supplies and with minimal change in temperature. We now investigate whether or not these periodic EBs can be used to compensate the oscillators in the presence of environmental variation. In this case, we investigate supply variation, which will be relevant to a battery-powered IoT node as its supply discharges.

To make a surface-level comparison to temperature, the supply variation of the RF oscillator is approximately 2.67 ppm/mV, and the temperature variation is approximately 40 ppm/°C. This means we can consider a supply change of 100 mV to emulate a temperature change of 6.7 °C.

The network synchronization timer (RFTimer) in Fig. 6, is compensated during supply changes using the feedback loop shown in Fig. 7. This feedback loop is unusual because the measurement is rarely sampled, and has a high level of non-white noise. After experimentation, a proportional controller was used with hysteresis about the target value guard time in an attempt to dampen the effects of high-noise transients, as summarized in the following expression, where k is set depending on the EB rate in the network (it is 100 μ s for a 4 s EB rate, and 500 μ s for a 16 s EB rate).

$$f = \begin{cases} 0 & |t_{ERR}| \leq k \\ -k/2 & t_{ERR} > k \\ +k/2 & t_{ERR} < -k \end{cases}$$

The other on-chip oscillators that were compensated were the RF oscillator (channel), the receiver sampling clock, and the transmit chipping clock. The channel frequency was maintained by observing the downconverted intermediate frequency. The sampling clock was kept constant by comparing its count value to the length of the incoming packet. A summary of these results with a 2.5 mV/s (150 mV/minute) supply ramp is shown in Fig. 8. To put this in perspective, this would be comparable to the discharge of a 56 mF capacitor over the course of two minutes with the chip’s “sleep” current of 140 μ A, or comparable to the discharge of a 100 μ Ah battery over approximately 15 minutes.

V. CONCLUSIONS

We have presented and demonstrated a multi-hop mesh network that uses a single crystal-enabled transceiver and four crystal-free transceivers. The network is capable of surviving variations in environment (in this case, variations in supply voltage) and the network’s ability to synchronize is limited exclusively by the jitter of the on-chip oscillators. The apparently constant relative chip-to-chip jitter of the crystal free transceivers allows for a fixed guard time throughout the network which simplifies scheduling.

The number of hops is limited by the number of devices that were available for experiment. In the future, more hops should be added to demonstrate that the node-to-node time error does indeed remain constant regardless of how deep in the network the link is. It is also not immediately clear whether

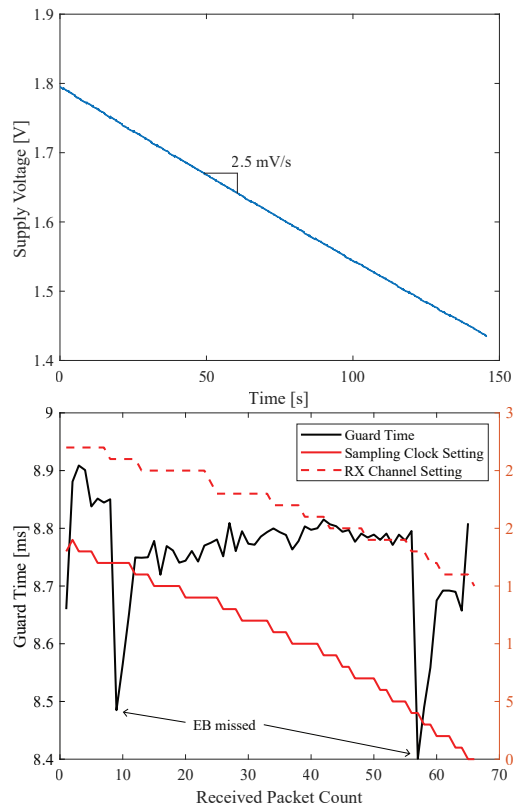


Fig. 8. On-chip oscillator calibration using received beacons from a crystal-enabled transceiver. The 2.5 mV/s ramp is approximately equivalent to a variation of 10.3°C per minute. Guard time is set unnecessarily high to increase packet reception likelihood.

the proposed protocol scales well with large numbers of nodes, or in star networks as opposed to multi-hop networks.

In addition, even under ambient conditions, the packet delivery rate (PDR) for crystal-free to crystal-free links was only around 90%. It is still unclear whether the PDR is low because of low SNR, incorrectly optimized transmit and receive settings, or channel conditions such as multipath fading.

Finally, in the experiments in this work, time synchronization was performed at the absolute worst time: as far after the beacon from a node's parent as possible. An interesting topic for future research is the over-the-air time compensation loop shown in Fig. 7. It is a fairly unusual control system, and further optimization could improve the minimum guard times presented in this work.

ACKNOWLEDGMENT

The authors would like to thank TSMC for donating the silicon used for the networking experiments, and Intel Corporation for funding during the design of the chip. This project has received funding from the European Union's Horizon Europe Framework Programme under Grant Agreement No. 101093046.

REFERENCES

[1] "IEEE Standard for Local and Metropolitan Area Networks, Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)," *IEEE Std*

802.15.4-2011 (Revision of IEEE Std 802.15.4-2006), pp. 1–314, Sep. 2011.

[2] B. SIG, "Bluetooth Specification, Core System Package, Low Energy Controller Volume," *Bluetooth Core Specification v5.0*, pp. 2533–2790, Dec. 2016.

[3] SiTime, "SiT1532 Smallest Footprint (1.2 mm²) CSP 10 ppm Ultra-Low Power 32.768 kHz XTAL Replacement rev. 1.28," Santa Clara, CA, USA, Mar. 2021, pp. 1–12. [Online]. Available: <https://www.sitime.com/products/32-khz-oscillators/sit1532>

[4] F.-W. Kuo, S. Pourmousavian, T. Siriburanon, R. Chen, L.-C. Cho, C.-P. Jou, F.-L. Hsueh, and R. B. Staszewski, "A 0.5V 1.6mW 2.4GHz Fractional-N All-Digital PLL for Bluetooth LE with PVT-insensitive TDC sing Switched-Capacitor Doubler in 28nm CMOS," in *2017 Symposium on VLSI Circuits*, 2017, pp. C178–C179.

[5] F. Maksimovic, B. Wheeler, D. C. Burnett, O. Khan, S. Mesri, I. Suci, L. Lee, A. Moreno, A. Sundararajan, B. Zhou *et al.*, "A Crystal-Free Single-Chip Micro Mote with Integrated 802.15.4 Compatible Transceiver, sub-mW BLE Compatible Beacon Transmitter, and Cortex M0," in *2019 Symposium on VLSI Circuits*. IEEE, 2019, pp. C88–C89.

[6] T. Watteyne, X. Vilajosana, B. Kerkez, F. Chraim, K. Weekly, Q. Wang, S. Glaser, and K. Pister, "OpenWSN: a Standards-Based Low-Power Wireless Development Environment," *Transactions on Emerging Telecommunications Technologies*, vol. 23, no. 5, pp. 480–493, 2012. [Online]. Available: <http://dx.doi.org/10.1002/ett.2558>

[7] K. S. Kim, S. Lee, and E. G. Lim, "Energy-Efficient Time Synchronization Based on Asynchronous Source Clock Frequency Recovery and Reverse Two-Way Message Exchanges in Wireless Sensor Networks," *IEEE Transactions on Communications*, vol. 65, no. 1, pp. 347–359, 2017.

[8] S. Ping, "Delay Measurement Time Synchronization for Wireless Sensor Networks," 2003.

[9] W. Kester, "Converting Oscillator Phase Noise to Time Jitter," *Analog Devices Tutorial MT-008 Rev. A*, p. 10, Oct. 2008.

[10] Keysight Technologies, "Phase Noise Measurement Application Measurement Guide," *Manual Part Number N9068-90002*, pp. 1–51, Feb. 2016. [Online]. Available: literature.cdn.keysight.com/litweb/pdf/N9068-90002.pdf

[11] T. Chang, T. Watteyne, F. Maksimovic, B. Wheeler, D. C. Burnett, T. Yuan, X. Vilajosana, and K. S. J. Pister, "QuickCal: Assisted Calibration for Crystal-Free Micromotes," *IEEE Internet of Things Journal*, vol. 8, no. 3, pp. 1846–1858, 2020.

[12] X. Huan, K. S. Kim, S. Lee, E. G. Lim, and A. Marshall, "A Beaconless Asymmetric Energy-Efficient Time Synchronization Scheme for Resource-Constrained Multi-Hop Wireless Sensor Networks," *IEEE Transactions on Communications*, vol. 68, no. 3, pp. 1716–1730, Mar. 2020, conference Name: IEEE Transactions on Communications.

[13] —, "Improving Multi-Hop Time Synchronization Performance in Wireless Sensor Networks Based on Packet-Relaying Gateways With Per-Hop Delay Compensation," *IEEE Transactions on Communications*, vol. 69, no. 9, pp. 6093–6105, Sep. 2021, conference Name: IEEE Transactions on Communications.

[14] W. Jingchao, Z. Ruohan, and G. Weiwen, "Time Synchronization in Networks: A Survey," in *Proceedings of the 2nd International Conference on Control and Computer Vision*, ser. ICCCV 2019. New York, NY, USA: Association for Computing Machinery, Jun. 2019, pp. 121–126. [Online]. Available: <https://doi.org/10.1145/3341016.3341038>