

One-hop Out-of-band Control Planes for Multi-hop Wireless Sensor Networks

CHAOJIE GU, Nanyang Technological University, Singapore

RUI TAN, Nanyang Technological University, Singapore

XIN LOU, Illinois at Singapore, Singapore

Separation of Control and Data Planes (SCDP) is a desirable paradigm for low-power multi-hop wireless sensor networks requiring high network performance and manageability. Existing SCDP networks generally adopt an *in-band* control plane scheme in that the control-plane messages are delivered by their data-plane networks. The physical coupling of the two planes may lead to undesirable consequences. Recently, multi-radio platforms (e.g., TI CC1350 and OpenMote B) are increasingly available, which make the physical separation of the control and data planes possible. To advance the network architecture design, we propose to leverage on the long-range communication capability of the Low-Power Wide-Area Network (LPWAN) radios to form one-hop out-of-band control planes. LoRaWAN, an open, inexpensive, and ISM band based LPWAN radio, is chosen to prototype our out-of-band control plane called LoRaCP. Several characteristics of LoRaWAN such as downlink-uplink asymmetry and primitive ALOHA media access control need to be dealt with to achieve high reliability and efficiency. To address these challenges, a TDMA-based multi-channel transmission control is designed, which features an urgent channel and negative acknowledgment. On a testbed of 16 nodes, LoRaCP is applied to physically separate the control-plane network of the Collection Tree Protocol (CTP) from its Zigbee-based data-plane network. Extensive experiments show that LoRaCP increases CTP's packet delivery ratio from 65% to 80% in the presence of external interference, while consuming a per-node average radio power of 2.97 mW only.

CCS Concepts: • **Computer systems organization** → **Sensor networks**; • **Networks** → *Network manageability*;

Additional Key Words and Phrases: Wireless sensor networks, low-power wide-area networks, LoRaWAN, network manageability

ACM Reference Format:

Chaojie Gu, Rui Tan, and Xin Lou. 2019. One-hop Out-of-band Control Planes for Multi-hop Wireless Sensor Networks. *ACM Trans. Sensor Netw.* 1, 1, Article 1 (January 2019), 30 pages. <https://doi.org/0>

A preliminary version of this work appeared in The 37th Annual IEEE International Conference on Computer Communications (INFOCOM 2018). The research was supported in part by a CoE Seed Grant and a Start-up Grant at Nanyang Technological University.

Authors' addresses: Chaojie Gu, Nanyang Technological University, School of Computer Science and Engineering, N4-B02A-01, 50 Nanyang Avenue, Singapore, 639798, gucj@ntu.edu.sg; Rui Tan, Nanyang Technological University, School of Computer Science and Engineering, N4-02C-85, 50 Nanyang Avenue, Singapore, 639798, tanrui@ntu.edu.sg; Xin Lou, Illinois at Singapore, Advanced Digital Sciences Center, 1 Create Way, #14-02 Create Tower, Singapore, 138602, lou.xin@adsc-create.edu.sg.

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

© 2009 Association for Computing Machinery.

1550-4859/2019/1-ART1 \$15.00

<https://doi.org/0>

1 INTRODUCTION

By 2018, 35 billion smart objects will be deployed, forming *things networks* that are interconnected by Internet of Things (IoT) [1]. Many of these networks will follow the paradigm of multi-hop wireless sensor networks. For instance, wireless meshes are increasingly adopted to interconnect surveillance cameras [60] and vehicles [59]. Wireless sensors have been widely deployed for sensing and control of building environment and energy use. Bluetooth Low Energy (BLE) will support mesh networking soon [7]. Wireless connectivity is also critical to the vision of Industry 4.0. Utility and manufacturing systems are increasingly adopting wireless metering and monitoring [37].

The main advantage of Multi-Hop Wireless Sensor Networks (MH-WSNs) is that, during the deployment phase, a network can easily scale up to cover a large geographic area. A primary design principle for MH-WSNs is the use of distributed protocols (e.g., routing [20]), where each node independently performs various networking functions (e.g., data forwarding) based on local information. Thus, the *control plane* (i.e., determination of how to handle packets) and the *data plane* (i.e., carrying out control-plane decisions) of these distributed protocols are jointly implemented at each network node. However, a distributed scheme without the global view often yields suboptimal performance. Moreover, although the distributed scheme may work satisfactorily most of the time thanks to a decade of research, it is often complex and difficult to manage once the network is deployed.

To improve the network performance and manageability, some MH-WSNs, especially those deployed for mission-critical tasks, have adopted centralized network controls. For instance, WirelessHART, an MH-WSN standard that has been adopted in over 8,000 manufacturing systems [37], prescribes centralized routing control based on a global view of the network. It thus better achieves certain performance objectives (e.g., firm/soft real-time packet delivery). Similarly, ISA100.11a, another industry-oriented MH-WSN standard, also adopts centralized routing control and network management. For the routing in these MH-WSNs, the control plane is separated from the data plane, in that the routing control is implemented at a centralized node whereas other network nodes follow the routing schedule to forward data packets. However, all these MH-WSNs adopt *in-band* control planes, i.e., the control-plane messages such as network status reports and routing schedules are delivered by the data-plane networks.

The physical coupling between the control and data planes in the in-band scheme may lead to undesirable consequences. The wireless data-plane network is susceptible to external interference. Deteriorated data-plane links may lead to delayed deliveries or even losses of the control-plane messages, making the network less responsive to data-plane link quality variations. Moreover, when the data plane loses key routing nodes (e.g., due to node hardware/software fault and depletion of battery) or the control plane makes wrong control decisions (e.g., due to design defects or erroneous human operations), the data-plane network may fall apart to disconnected partitions. As a result, restorative network control commands in the control plane may not be able to reach the destination nodes. Recent research has studied protecting the control plane from data-plane faults [23]. However, the solution has limited protection capability against a single link failure only [23].

In light of the in-band scheme's pitfalls, this paper studies an *out-of-band* scheme, where the control plane uses a dedicated wireless network different from the data-plane network. The increasingly adopted multihoming and increasingly available dual-band or multi-band radios prepare the IoT hardware platforms for implementing the out-of-band scheme. The prevailing IoT platforms are generally equipped with multiple heterogeneous network interfaces: Raspberry Pi 3 supports Ethernet, Wi-Fi, and BLE; Firestorm [5] supports BLE and Zigbee; Arduino has various add-on boards to support different radios. Some latest IoT platforms are equipped with both short-range

and long-range low-power radios. The OpenMote B platform [48] integrates a Texas Instruments (TI) CC2538 Zigbee radio and an Atmel AT86RF215 sub-GHz long-range radio. The LoPy4 platform [51] offers both a low-power wide-area networking radio (Sigfox or LoRaWAN) and a short-range radio (Zigbee or Wi-Fi). In particular, several recent System-on-a-Chip (SoC) modules integrate both short-range and long-range low-power radios. For example, the TI CC1350 [25] provides both BLE and sub-GHz long-range radio. Note that CC1350 is one of several SoC modules adopted by the SensorTag platform [26]. The TI CC1352R [24], which is in pre-production phase at present, provides both Zigbee and sub-GHz long-range radios. Given the increasing availability of heterogeneous low-power radios, this paper aims to study the network performance and manageability improvements as well as the related overhead by using different radios to form physically separated data- and control-plane networks.

To design the out-of-band control-plane network for MH-WSNs, Zigbee, Wi-Fi Direct, and the coming BLE mesh are ill-suited, since otherwise the control-plane network will be yet another multi-hop network that suffers the same manageability and fragility issues as the data-plane network. Cellular networks provide pervasive connectivity. However, the cellular network radios consume excessive power. As measured in [30], the Long-Term Evolution (LTE) user equipment consumes about 2 W for both downlink and uplink transmissions. Instead, we propose to use the emerging Low-Power Wide-Area Network (LPWAN) technologies (e.g., LoRaWAN, Sigfox, Weightless, and other sub-GHz wireless) for the out-of-band control plane. From our measurements in this paper, two LPWAN technologies have a power consumption of about 0.1 W only. Moreover, owing to the kilometers communication range of LPWAN links, the LPWAN-based control plane can be a one-hop star network, greatly simplifying its deployment and management.

This paper investigates the suitability of two LPWAN technologies for designing out-of-band control planes for MH-WSNs, i.e., LoRaWAN and TI's sub-GHz long-range radio. These two LPWAN technologies use license-free ISM bands and do not rely on managed infrastructures. Extensive measurements are conducted to profile the power consumption, timing performance, and indoor communication performance of the two radios. While both of them can meet the technical requirements for building control planes, this paper focuses on using LoRaWAN to prototype the proposed one-hop control plane and gain insights. While the low-power long-range communication capability is the key advantage of LoRaWAN, the following three limiting characteristics of LoRaWAN need to be managed properly. First, a LoRaWAN downlink frame from the controller to a network node must be in response to a precedent uplink frame. Thus, the transmissions of network control commands initiated by the controllers may be postponed to the network node's status reporting. Second, LoRaWAN supports uplink concurrency but no downlink concurrency. This downlink-uplink asymmetry impedes acknowledging each uplink frame, whereas the control plane generally desires reliable message delivery. Third, LoRaWAN adopts the ALOHA Media Access Control (MAC) protocol, which may perform poorly in traffic surges.

To address these issues, this paper presents the design and implementation of a prototype system called *LoRaCP* (long-range control plane). Based on our extensive measurements on LoRaWAN's energy and latency profiles, we design *LoRaCP-TxC*, a TDMA-based multi-channel transmission control approach featuring uplink heartbeats, the negative acknowledgment (NAK), and an ALOHA-based urgent channel, to manage the transmissions of the control-plane messages. The uplink heartbeats open downlink windows for controller-initiated network commands and maintain network nodes' clock synchronization for TDMA. With NAK, the controller does not need to acknowledge every uplink frame. The urgent channel complements the TDMA channels to mitigate the rigidity of TDMA. On a testbed of 16 nodes, LoRaCP is applied to physically separate

the control plane of the Collection Tree Protocol (CTP) [20] from its Zigbee-based data-plane network. The LoRaCP-TxC is implemented in the application layer using the program library of a LoRaWAN platform. Extensive experiments show that LoRaCP increases CTP's packet delivery ratio from 65% to 80% in the presence of external interference, while consuming a per-node average radio power of 0.9 mA only under an operating voltage of 3.3 V, much lower than the active power of many recent sensor platforms' microcontrollers (e.g., 8.6 mA on Firestorm [5]).

The main contributions of this paper are summarized as follows:

- A set of comparative TinyOS simulations are conducted to study the performance of CTP in an MH-WSN under distributed versus centralized network controls, and in-band and out-of-band centralized network controls.
- Measurement studies are performed to profile the power consumption, timing accuracy, and indoor communication performance of two LPWAN radios, i.e., LoRaWAN and TI's sub-GHz long-range radio.
- A LoRaWAN-based out-of-band control plane prototype system called LoRaCP is designed and implemented. Experiments on a testbed of 16 nodes show the network resilience brought by LoRaCP for CTP.

The rest of the paper is organized as follows. §2 reviews related work. §3 presents a number of simulation-based examples to motivate the out-of-band scheme for control planes. §4 profiles the performance and overhead of LoRaWAN and TI's sub-GHz long-range radio. Based on the profiling results, §5 designs LoRaCP. §6 presents the evaluation results of LoRaCP in testbed experiments. §7 concludes this paper.

2 RELATED WORK

The increasing availability of multihoming and multi-band radios enables researchers to investigate the benefits by leveraging on multiple radios. Existing studies that exploit multiple radios can be broadly divided into two classes of *bandwidth aggregation* and *Separation of Control and Data Planes* (SCDP).

Bandwidth aggregation uses multiple network interfaces to transmit/receive data simultaneously to increase throughput. Habak et al. [21] surveyed early bandwidth aggregation literature, and categorized them into solutions at the application, transport, network, and link layers. Early studies include the Multi-Radio Unification Protocol (MRUP) [4] and system designs of multi-radio approaches [6]. The MRUP proposed by Adya et al. [4] runs at the link layer to coordinate the operations of multiple wireless network cards tuned to non-overlapping frequency channels, based on locally available information only. Subsequently, Bahl et al. [6] study various design issues of the multi-radio approaches in the hardware, algorithmic, and protocol aspects. Recent development that is not covered by the survey paper [21] is reviewed as follows. These new studies are divided into two categories. The first category exploits *homogeneous* radios. FatVAP [31] enables a 802.11 wireless card to connect to multiple access points. FastForward [17] uses two 802.15.4 radios operating on different channels, with one receiving and the other forwarding data simultaneously. The second category exploits *heterogeneous* radios. In [39], various trade-offs in designing energy efficient multi-radio platforms are investigated. In MicroCast [33], a group of smartphones cooperate in downloading a video from the cellular and share their downloaded portions through device-to-device links (e.g., Wi-Fi). MultiNets [46] deals with the switching between multiple network interfaces on mobile devices. In [44], Mu et al. optimize the selection of radios and their transmission powers. Recent studies [36, 45] characterize the performance and energy consumption of Multipath TCP through multiple radios of a mobile device. Different from bandwidth aggregation

that combines multiple network interfaces in the data plane to increase throughput, SCDP aims to improve network optimality and manageability.

Software-Defined Networking (SDN), with SCDP as its core concept, is a growing momentum in data-intensive networks (e.g., data center and enterprise backbone networks). To avoid the undesirable coupling between the control and data planes, SDN recommends the out-of-band scheme [47]. SCDP can be naturally applied in wireless local area networks and cellular networks, as their topologically centralized access points and base stations can run the control-plane logics for better resource allocation and mobile node handover [29]. However, there is limited research on SCDP in multi-hop wireless networks. An OpenFlow-enabled Wi-Fi mesh was built in [14], where each Wi-Fi card is split into two virtual interfaces with different Service Set Identifiers (SSIDs) and the two planes are two multi-hop networks in their respective SSIDs. In the initial thinking of applying SCDP in wireless sensor networks [38], the design choice of in-band or out-of-band control plane is dubious. To the best of our knowledge, WASP [32] is the only system that implements out-of-band control plane for multi-hop wireless networks. WASP uses Wi-Fi Direct and cellular network of smartphones to form the data and control planes, respectively. Different from WASP, this paper focuses on low-power networks with a limited energy budget.

LoRaWAN, an emerging LPWAN technology, has received increasing research in recent years. Existing studies mainly aim to improve the communication performance and battery lifetime of the LoRaWAN end devices. LoRaWAN's communication performance is profiled via field measurements [40, 50, 57]. A recent work [35] provides an extensive and in-depth measurement study in various outdoor environments. The work [56] presents the design of a multi-channel and multi-hop MAC protocol for a LoRa-based wildlife monitoring system. The work [15] presents an approach that can predict LoRaWAN link quality based on multispectral images from remote sensing. The Choir [18] system exploits the diverse frequency biases of the LoRaWAN end devices to decode colliding frames from different end devices. The Charm [16] system exploits coherent combining to decode a frame from the weak signals received by multiple geographically distributed LoRaWAN gateways. In addition to the above studies on improving the scalability and robustness of LoRaWAN networks, several recent studies have proposed various backscatter designs for LoRa [22, 49, 55, 58]. Different from the off-the-shelf LoRa nodes that use their own power sources to drive the radio transmissions, a backscatter device generates backscatter signals in response to some activation signal to transmit bits. Thus, backscatter can extend the battery lifetime. Different from these studies, this paper focuses on exploiting LoRa's long-range communication capability to improve the data plane's network performance and resilience against external disturbances.

3 SIMULATION STUDIES

In this section, comparative simulation studies are conducted to motivate the use of out-of-band centralized network control to improve network performance and resilience. Specifically, §3.1 compares the distributed network control scheme and the centralized network control scheme. §3.2 compares the in-band centralized network control and the out-of-band centralized network control. All the simulations are conducted in the TinyOS simulator TOSSIM [34].

3.1 Distributed versus Centralized Network Control

This section compares through simulations the network performance achieved by the CTP [20] and its centralized variant that is called CTP-SCDP. In §6, LoRaCP will be applied to implement CTP-SCDP and evaluated on a testbed. This paper uses CTP as the case study network protocol, because CTP has an open implementation and is a standard component of the industry-class TinyOS Production operating system [2]. The results based on CTP will provide insights into the performance improvement by SCDP and showcase the use of LoRaCP to physically separate the control

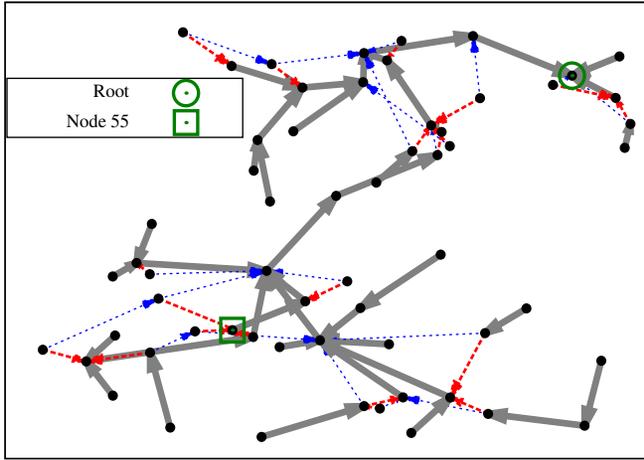


Fig. 1. Node placement and the routing trees constructed by CTP and CTP-SCDP. The solid thick gray links are shared by the CTP and CTP-SCDP trees; the dashed thick red links are on the CTP tree only; the dashed thin blue links are on the CTP-SCDP tree only.

and data planes. The obtained insights are also useful to the SCDP designs of other MH-WSN protocols.

CTP aims to maintain a minimum-cost routing tree in the presence of dynamic link quality characterized by the Expected Transmission Count (ETX). The cost of a route to the tree root is the sum of the ETXs of the links on the route. A node i estimates the route cost using the Residual ETX (RETX), which is given by $RETX_i = ETX_{i,p} + RETX_p$, where $ETX_{i,p}$ is the ETX of the link between node i and its parent node p , and $RETX_p$ is node p 's RETX. CTP works in a distributed manner, in that each node i selects its parent p from the set of its neighbor nodes \mathcal{N} based on local information only. Specifically, $p = \arg \min_{j \in \mathcal{N}} ETX_{i,j} + RETX_j$, where $ETX_{i,j}$ is estimated based on the transmissions of beacons and data frames; $RETX_j$ is broadcast in node j 's beacons.

In CTP, the information about the quality of a link propagates to the whole network during the beaconing process. However, this propagation takes time. Thus, when link quality changes over time, the RETX of any node i cannot capture the latest ETXs of the links on its route to the root. In particular, the closer the links on the route are to the root, node i 's knowledge about the links (which is encompassed in $RETX_i$) is more out-of-date. As a result, CTP may not construct the minimum-cost tree in the presence of time-varying link quality. Differently, in CTP-SCDP, the latest ETXs can be updated to the network controller in time. Specifically, when the change of any ETX of a node exceeds a certain threshold, the node can send the latest ETX to the network controller via the dedicated control-plane network. Upon receiving an updated ETX, the network controller shall recompute the optimal routing and send the changes in routing to concerned nodes via the control-plane network as well.

Simulations are conducted in TOSSIM to compare CTP and CTP-SCDP. A total of 60 nodes are placed randomly in a 200 m \times 200 m region as illustrated in Fig. 1. Note that the size of the simulated region is similar to the dimension of a building (190 m long) in which we deploy LoRaWAN and TI's sub-GHz nodes to measure LPWAN's communication performance in §4. The density of the simulated nodes is similar to the building environment monitoring applications, e.g., temperature distribution monitoring in data centers and hazard dust/gas concentration monitoring on factory floors. Link gains are generated according to the Euclidean distances between nodes using a tool

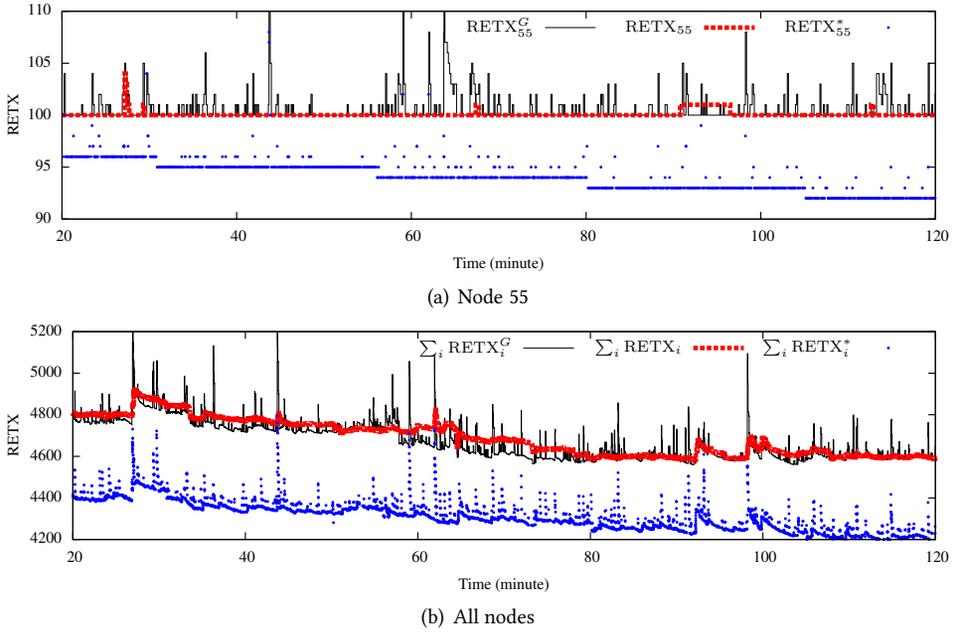


Fig. 2. The groundtruth RETX and estimated RETX in CTP as well as the optimal RETX if CTP-SCDP is adopted. The results show that CTP cannot build a minimum-cost tree.

in TOSSIM. TOSSIM applies a radio propagation model for each node. It can also simulate the RF noises and interference that a node is subjected to. Radios' hardware noise floor is set to be -90 dBm, which is a mild noise level. To simulate CTP-SCDP, a node is added as the network controller, which has sufficiently large link gains with any other nodes, such that the control-plane network is a one-hop star network. The TOSSIM is configured such that the data-plane links do not interfere with the control-plane links. As this one-hop star network is not used to transfer sensor data, this CTP-SCDP system follows the out-of-band control plane scheme. In CTP-SCDP, node i sends the latest $ETX_{i,j}$ to the network controller. Upon receiving an ETX update, the controller updates a directed graph with the ETXes as the edge costs and recomputes the minimum-cost routing tree using the Dijkstra's algorithm. Then, the controller sends the new parent information to the nodes.

Two sets of simulations are conducted to show the benefits of SCDP. The first set shows the suboptimal performance of CTP. Specifically, CTP and CTP-SCDP run concurrently, but the controller in CTP-SCDP does not send routing control commands to the nodes. Thus, the routing is managed by CTP only. The following evaluation metrics are considered:

- (1) RETX of node i estimated by CTP (denoted by $RETX_i$) and the sum of all RETXes (denoted by $\sum_i RETX_i$);
- (2) The *ground-truth* RETX of the route determined by CTP for node i (denoted by $RETX_i^G$), which can be measured as the sum of the latest ETXes of the links on the route obtained by the controller in CTP-SCDP, as well as the sum of all ground-truth RETXes (i.e., $\sum_i RETX_i^G$);
- (3) The minimum RETX of node i computed by CTP-SCDP (denoted by $RETX_i^*$) and the sum $\sum_i RETX_i^*$.

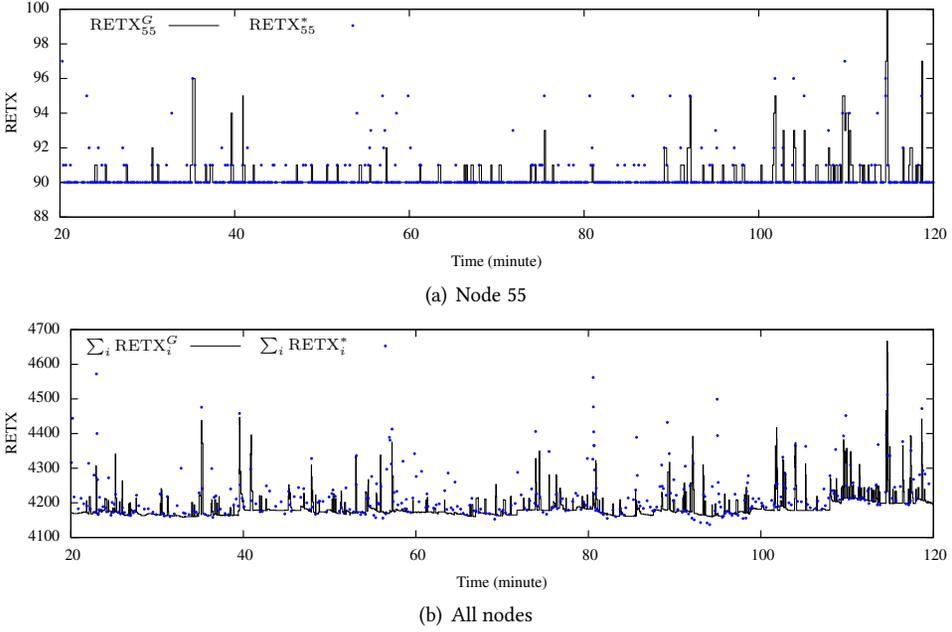


Fig. 3. The groundtruth RETX achieved by CTP-SCDP and the true optimal RETX. The results show that CTP-SCDP can build a minimum-cost tree.

The simulated time duration is two hours, during which each node generates a data packet every eight seconds. Fig. 1 shows the routing trees computed by CTP and CTP-SCDP at the end of the simulation. They are different. Fig. 2 shows the evaluation metrics for Node 55 and all the nodes over the two hours. It can be seen that, compared with the ground truth (i.e., the solid black curves), CTP’s knowledge about the chosen routes (i.e., the dashed red curves) cannot capture many transient changes in the ground truth, because of the information propagation latency in the distributed network control. Compared with the global optimal (i.e., the blue dots), the routes chosen by CTP have higher costs.

In the second set of experiments, only CTP-SCDP runs. Fig. 3 shows the results. It can be seen from the figure that the routes chosen by CTP-SCDP generally achieve the minimum costs. The above two sets of simulations show that the centralized network control improves the network performance in dynamic network conditions. Thus, the centralized control enabled by SCDP is desirable for performance-critical networks such as those deployed for industrial applications [37].

3.2 In-band versus Out-of-Band Network Control

Our simulations in §3.1 demonstrate the underperformance of distributed network control. As discussed in §1, a number of mission-critical MH-WSNs have adopted centralized network control to improve network performance and manageability. They all follow the in-band control plane scheme. However, the physical coupling of the control and data planes generates various challenges. For instance, given the fragile nature of wireless, how to protect the in-band control plane against data-plane link failures is a challenging problem. Recent research has investigated this issue. Nevertheless, existing solutions provide limited protection capability. For instance, the solution proposed in [23], though sophisticated, can handle a single link failure only. The in-band

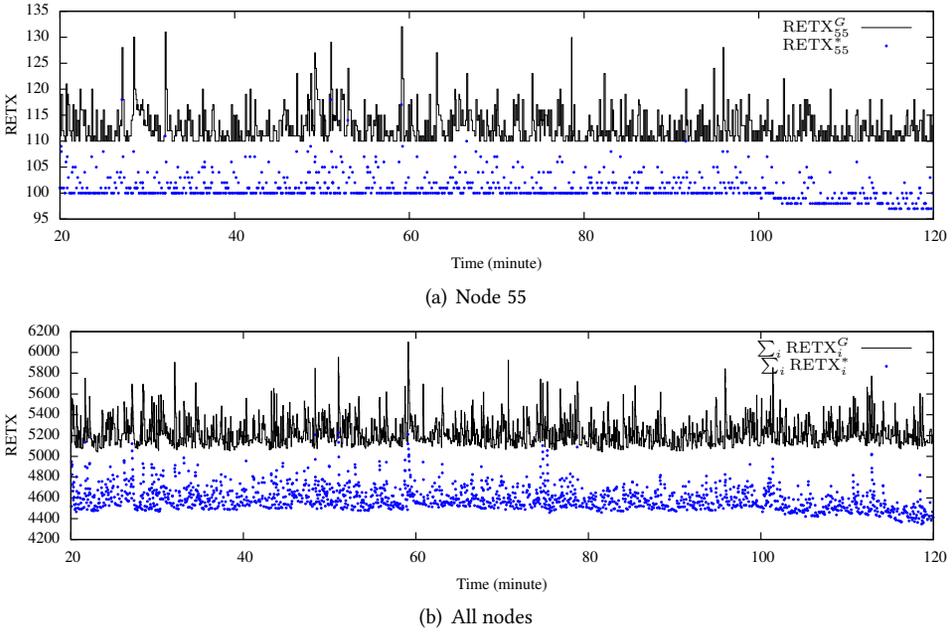


Fig. 4. The groundtruth RETX achieved by the in-band CTP-SCDP and the true optimal RETX. The results show that the in-band CTP cannot build a minimum-cost tree.

control plane protection under a general setting is still an open issue. In this section, simulations are conducted to compare the CTP systems with in-band and out-of-band centralized network controls, which are referred to as in-band CTP-SCDP and out-of-band CTP-SCDP, respectively. The out-of-band CTP-SCDP in this section is same as the CTP-SCDP in §3.1. In the in-band CTP-SCDP, each node will report the link quality information to the network controller using the multi-hop data-plane network. Upon receiving the link quality information from a node, the controller will compute the optimal routing path for this node. Then, the controller disseminates the new routes using targeted messages to the individual nodes. In both the in-band and out-of-band CTP-SCDP, a packet will be dropped after 30 unsuccessful (re-)transmissions to the next hop. In what follows, the in-band and out-of-band schemes are compared in terms of resilience to link noises and node faults, as well as the network convergence speed.

3.2.1 Resilience to link noises. The simulation settings in §3.1 are adopted in this section. The radios' hardware noise floor is set to be -90 dBm. The simulated time duration is two hours, during which each node generates a data packet every eight seconds. Fig. 4(a) and Fig. 4(b) show the evaluation metrics of the in-band control scheme for Node 55 and all the nodes over the two hours. The results show that the in-band control scheme cannot achieve optimal network performance. Compared with the true optimal network performance that is achieved by the out-of-band CTP-SCDP (i.e., the blue dots), the routes chosen by the in-band CTP-SCDP have higher cost. This is because, with the in-band CTP-SCDP, the data-plane and control-plane packets delivered by the same data-plane network may collide. The delayed and unsuccessful deliveries of the control-plane packets will lead to performance degradation of the data-plane network. Due to the limited payload size of a TinyOS message, the need of sending new routes to different nodes in separate packets also increases the contention between the data-plane and control-plane networks. Differently, in

Table 1. The packet delivery ratios (PDRs) achieved by the in-band and out-of-band CTP-SCDP schemes under different settings of radio noise floor.

Noise floor setting (dBm)	-105	-100	-95	-90	-85
PDR of in-band CTP-SCDP	99.97%	99.87%	99.80%	99.80%	8.48%
PDR of out-of-band CTP-SCDP	100%	100%	100%	100%	10.90%

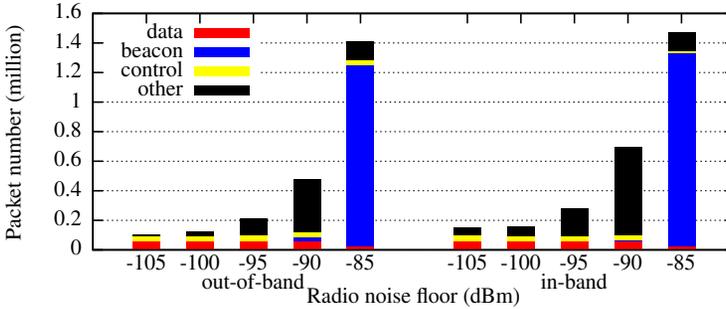


Fig. 5. The breakdown of the total transmitted packets in the out-of-band (left group) and in-band (right group) CTP-SCDP networks under various radio noise floor settings. The bars labeled “data”, “beacon”, “control” represent the data packets, beacons, and control-plane packets generated by all the nodes; the bar labeled “other” represents the packets (either data or control packets) successfully forwarded and re-transmitted packets by all the nodes.

the out-of-band CTP-SCDP, the data-plane and control-plane packets will not interfere with each other.

The data packet delivery ratios (PDRs) achieved by the in-band and out-of-band CTP-SCDP networks under various settings of the radio noise floor are measured. Table 1 shows the results. It can be seen that when the noise floor is from -105 dBm to -90 dBm, both the in-band and out-of-band CTP-SCDP networks can deliver (almost) all data packets to the sink. When the noise floor setting is -85 dBm, both networks have similarly low PDRs. This is because the link quality of the data-plane network is too poor to support reliable packet transmissions.

The numbers of packets in the following four categories of the whole network are counted to better understand the differences between the in-band and out-of-band schemes: (1) data packets generated by the nodes, (2) beacons, (3) control packets generated by the nodes, and (4) others including the packets (either data or control packets) successfully forwarded and re-transmitted by the nodes. Fig. 5 shows the above counts in the in-band and out-of-band CTP-SCDP networks under various radio noise floor settings. We now summarize the observations from Fig. 5 and discuss the reasons of these observations.

- When the noise floor setting increases from -105 dBm to -90 dBm, the increases of re-transmissions are observed. This is due to the deteriorated link quality. Moreover, the in-band CTP-SCDP network always has more re-transmissions than the out-of-band CTP-SCDP network. This is because the data and control planes of the in-band CTP-SCDP network contend for the bandwidth. The delayed or unsuccessful deliveries of the control-plane packets also lead to the sub-optimality of the data collection tree on the data-plane network as observed earlier. The sub-optimality leads to increased total traffic in return, although the in-band CTP-SCDP network maintains nearly 100% PDRs as shown in Table 5.

- When the noise floor setting is -85 dBm, the number of successful re-transmissions is reduced compared with that when the noise floor is -90 dBm. This is because the networks have excessive dropped packets due to the poor link quality. Moreover, the nodes transmit beacons intensively to track the highly dynamic link quality. Under this noise floor setting, from Table 1, both the in-band and out-of-band CTP-SCDP networks have low PDRs. Compared with the results for lower noise floor settings, the data packets both in the out-of-band and in-band CTP-SCDP networks reduce. This is because there are too many data packages to be re-transmitted and only a portion of generated data packets are processed by the end of the simulation.

It can be seen from the above results that, when the network can maintain good PDRs in the presence of link noises, the out-of-band scheme can reduce the cost in forwarding and re-transmitting packets.

3.2.2 Network convergence. With a one-hop out-of-band control-plane network, the nodes' local information can be directly transmitted to the controller. In contrast, in the distributed scheme or the in-band SCDP network, the local information propagates to the controller hop by hop to the controller, resulting in longer delays. Therefore, the out-of-band CTP-SCDP can converge to a new routing schedule faster when the network boots or has changes (say due to nodes' movements). A set of experiments are conducted to compare the convergence of the networks under different network control schemes. A metric $D(t_0, t_1)$ is computed to characterize the convergence of the network within a time duration of (t_0, t_1) . Specifically, the convergence metric is given by $D(t_0, t_1) = \frac{\sum_{t=t_0}^{t_1} (\sum_i \text{RETX}_i^G(t) - \sum_i \text{RETX}_i^*(t))^2}{t_1 - t_0}$, where $\text{RETX}_i^G(t)$ and $\text{RETX}_i^*(t)$ denote the ground-truth RETX of node i at time t and the minimum RETX of node i at time t , respectively. A smaller $D(t_0, t_1)$ value suggests a better convergence to the minimum-cost tree.

Our CTP-SCDP networks adopt the CTP's adaptive beaconing to keep track of the $\text{ETX}_{i,j}$ (i.e., the ETX between node i and node j). The network convergence highly depends on the frequency of the adaptive beaconing. We now briefly introduce the adaptive beaconing mechanism. More details can be found in [20]. When the network is booted or the CTP detects a potential change of ETX based on the (re-)transmissions of the data packets, CTP sends beacons every T_{bl} seconds, where T_{bl} is the lower bound of the beaconing interval. After that, it increases the beaconing interval exponentially up to T_{bu} , where T_{bu} is the upper bound. This adaptive beaconing saves radio energy when the link quality is stable. In CTP, the default settings for T_{bl} and T_{bu} are 128 ms and 5.12×10^5 ms, respectively. The first six error bars in Fig. 6 show the measured convergence metric $D(15 \text{ min}, 120 \text{ min})$ of the CTP when the T_{bu} increases from 5.12×10^2 ms to 5.12×10^7 ms and T_{bl} adopts its default setting. Specifically, each error bar shows the distribution of the convergence metric in a total of fifteen 2-hour simulations. It can be seen that the knee point of the average D value is at CTP's default setting for T_{bu} (i.e., 5.12×10^5 ms). With smaller settings for T_{bu} , the CTP network does not converge well (i.e., large D values). This is because, in CTP, both data packets and control packets have the same priority and the highly frequent beacons can result in contentin. With larger settings for T_{bu} , the distributions of the D value are more spread. This is because the CTP network updates the ETXs using less frequent beacons and the network may not react to network condition changes in time. Thus, the default setting of $T_{bu} = 5.12 \times 10^5$ strikes a satisfactory trade-off between the overhead of beaconing and performance in tracking ETXs timely for maintaining low cost.

The last two error bars in Fig. 6 show the results for the in-band and out-of-band CTP-SCDP when the default setting for T_{bu} is applied. It can be seen that the in-band CTP-SCDP network has worse network convergence compared with the CTP adopting the default T_{bu} setting. This is

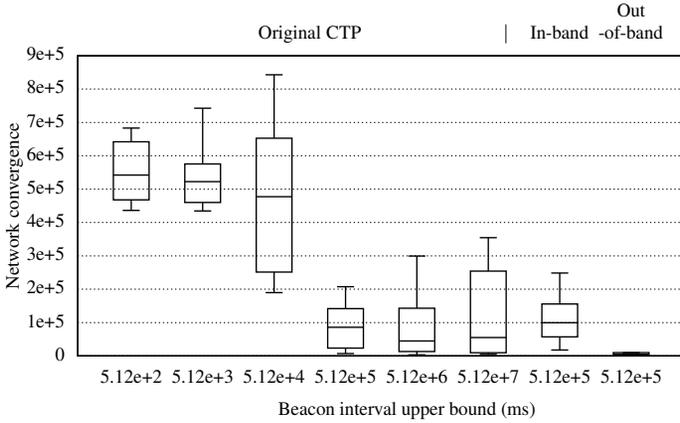


Fig. 6. Measured network convergence metric $D(15\text{th minute}, 120\text{th minute})$ of the original CTP, in-band CTP-SCDP, and out-of-band CTP-SCDP, under various settings of the beaconing interval upper bound.

because, compared with the distributed control scheme, the in-band scheme imposes additional overhead of conveying control packets to/from the centralized network controller using the data-plane network. The out-of-band CTP-SCDP achieves the best network convergence.

3.2.3 Resilience to node faults. In real deployments, WSN node faults (due to say battery depletion or hardware malfunction) are not uncommon. The distributed network control in general can well handle node faults. For instance, in CTP, the ETXs associated with a faulty node will quickly increase to infinity. As a result, the faulty node will not be selected as the parent node in the route to the sink. Differently, in a pure SCDP network, the decision of switching the parent node shall be made by the centralized network controller. If the SCDP network adopts the in-band scheme, the deliveries of the network status updates and parent node switching decisions may be affected by the node faults. For instance, in the simulations of the in-band CTP-SCDP, we switch off Node 55 shown in Fig. 1 in the 80th minute to simulate a node fault event. As a result, the nodes on the sub-tree rooted at Node 55 are disconnected from the network – they cannot send/receive data/control packets to/from the network controller. The network controller can only infer the occurrence of the fault based on the increasing ETXs reported by the neighbors of Node 55 that are still connected with the network controller. To stick to the pure SCDP scheme with in-band control plane, a separate network restoration mechanism will be needed. In particular, the control packets to the disconnected sub-tree cannot be delivered using the old routes. A message flooding may be needed to deliver the control packets. Differently, with an one-hop out-of-band control-plane network, since every node is directly connected with the network controller, new routes can be sent to the affected nodes to isolate the faulty node.

3.3 Summary of Simulation Results

Our performance profiling for CTP and its various variants shows that (i) the centralized control can further reduce the cost of the data collection tree compared with the original distributed control scheme (§3.1); (ii) the out-of-band SCDP network outperforms the in-band SCDP network in terms of network convergence as well as resilience to link noises and node faults (§3.2). These results suggest that the out-of-band SCDP scheme is a promising design for the MH-WSNs that have certain performance and resilience requirements. The rest of this paper will study various aspects of the out-of-band SCDP scheme, which include its real implementation using LPWAN

technologies, additional energy consumption, and the resulted data-plane network performance improvement under real-world settings.

4 LORAWAN AND TI'S SUB-GHZ PERFORMANCE PROFILING

This section profiles the energy consumption and latency of two LPWAN technologies, i.e., LoRaWAN and TI's sub-GHz, which are important to the design of the one-hop out-of-band control planes.

4.1 LoRaWAN and Its Characteristics

4.1.1 Introduction of LoRaWAN. LoRaWAN is an open data link layer specification based on LoRa, a proprietary PHY layer technique that uses a Chirp Spread Spectrum modulation and operates in sub-GHz ISM bands (e.g., EU868 MHz and US915 MHz). LoRa admits configuring the ratio between the symbol rate and chip rate by specifying an integer Spreading Factor (SF) within [6, 12]. Specifically, each symbol is modulated by 2^{SF} chips. A higher SF increases the signal-to-noise ratio and the communication range, but decreases the symbol rate. In this paper, six SF settings (from SF7 to SF12) are used.¹ The communications using different SFs are orthogonal and thus can be concurrent. Thus, in this paper, the terms *SF* and *channel* are used interchangeably. LoRa also admits configuring bandwidth and coding rate. These two parameters can affect the communication performance of LoRa. In this paper, we configure them to be 125 kHz and 4/5, respectively. Our performance profiling can be easily extended to address other settings of these two parameters.

A LoRaWAN network is formed by one or more *gateways* and many *end devices*. The gateway, often Internet-connected, can simultaneously handle the communications with multiple nodes in different channels. LoRaWAN defines three classes (A, B, and C) of end devices. A Class-A device's uplink transmission is followed by two downlink windows (RX1 and RX2). Downlink communications to the node at any other time will have to wait until the next uplink from the node. As Class-A is the most power efficient and supported by any end device, Class-A is chosen to design the out-hop out-of-band control plane.

4.1.2 Characteristics of LoRaWAN. The low-power long-range communication capability is the main advantage of LoRaWAN that makes it promising for control planes of MH-WSNs. However, we need to keep in mind the following two limiting characteristics of LoRaWAN in the design of the one-hop out-of-band control plane.

- **Downlink-uplink asymmetry:** LoRaWAN is mainly designed and optimized for uplinks from end devices to gateway. For instance, the LoRaWAN concentrator can receive frames from multiple channels simultaneously, whereas it can send a single downlink frame only at a time. Moreover, the Class-A specification requires that any downlink transmission must be unicast, in response to a precedent uplink transmission.
- **Lossy links:** From existing tests [40], with SF12, the frame reception rate is about 80% at a distances of 2.5 km. To build a reliable control-plane network, the frame losses need to be dealt with properly. Acknowledging each uplink frame is wasteful given the scarce downlink time as discussed earlier.

Moreover, the following two default features of LoRaWAN need to be considered and/or re-engineered in the design of an efficient one-hop out-of-band control plane.

- **ALOHA MAC:** LoRaWAN uses ALOHA that may perform poorly in surges of control plane messages. Moreover, as LoRa does not prescribe carrier sense capability, CSMA is not viable. Time-Division Multiple Access (TDMA) is often adopted for reliability that control planes

¹According to a LoRa chip's datasheet [54], SF6 is a special setting that is not enabled by default. Thus, we do not use SF6.

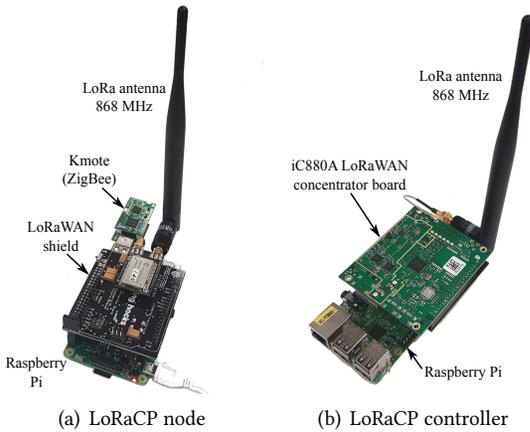


Fig. 7. LoRaCP hardware prototypes. (The Raspberry Pi is for fast prototyping only; it will not be needed if LoRa is built into the MH-WSN platform.)

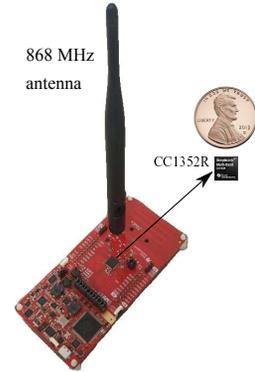


Fig. 8. CC1352R launchpad used in our measurement study. The launchpad includes various peripherals that will not be needed in real deployments. The CC1352R consists of MCU, Zigbee and sub-GHz radios.

desire. However, as shown in this paper, the implementation of TDMA on LoRaWAN is non-trivial. Moreover, a strict TDMA may result in undesirable delays in transmitting urgent messages. In addition, for the universality of LoRaCP, a transmission control protocol that does not need to modify LoRaWAN’s MAC-layer code is desired.

- **Text transmission only:** LoRa admits hexadecimal ASCII string only. To send an integer, it transmits the hexadecimal ASCII code word of each literal character of the integer. For instance, LoRa transmits “3130303030” (i.e., a character string of ten bytes) to represent a two-byte integer 10000. Thus, LoRa is inefficient in transmitting large integers (e.g., an epoch timestamp in ms resolution). However, a control plane will need to transmit many integers for control purpose.

In the design of the one-hop out-of-band control plane (cf. §5), the downlink-uplink asymmetry and lossy links will be managed by the NAK mechanism. Moreover, we will develop a new ASCII coding scheme for integers and design a TDMA-based multi-channel transmission control protocol on top of LoRaWAN’s ALOHA MAC. Both features can be implemented in the application layer.

4.2 LoRaWAN Performance Profiling

This section profiles the performance of LoRaWAN using our LoRaCP hardware prototypes. The results are important to the software design of LoRaCP in §5.

4.2.1 LoRaCP hardware prototypes. Performance profiling is performed based on the following prototype hardware platforms. Each end device integrates a Cooking Hacks LoRaWAN shield and a Raspberry Pi (RPi) 3 Model B single-board computer. The shield has a Microchip RN2483 LoRaWAN chip, an 868 MHz antenna, and interfacing circuits. The shield can be controlled by the RPi using a C++ library from Cooking Hacks. The gateway integrates an RPi and an IMST iC880A LoRaWAN concentrator board [3]. The iC880A board can receive frames over all LoRa channels simultaneously.

Table 2. Current consumption of a LoRaWAN 868 MHz module with bandwidth configured to 125 kHz.

SF	7	8	9	10	11	12
Transmitting current (mA)	39.741	39.861	39.495	39.694	40.441	40.413
Receiving current (mA)	14.196	14.208	14.227	14.240	14.234	14.240

A Zigbee-based Knode is plugged into a USB port of the RPi of each end device, forming a *LoRaCP node*. The nodes use their Zigbee radios to form the data-plane network. In this paper, we choose Zigbee because there is a body of MH-WSN protocols implemented for Zigbee. From now on, the gateway is referred to as *LoRaCP controller*. The controller unnecessarily has a Zigbee radio, since it may not be in the data-plane network. RPi is used to quickly prototype the integration of LoRaWAN and Zigbee. The results of this paper will suggest that integrating LoRaWAN into the design of MH-WSN platforms, especially those desiring high network performance and manageability, is valuable. In such designs, the RPi will not be needed. Fig. 7 shows our prototypes.

4.2.2 Power and energy profiling. From RN2483's datasheet, its current consumption during transmitting, receiving, and sleeping modes with a supply voltage of 3.3 V are 38.9 mA, 14.2 mA, and 0.0016 mA, respectively [41, 42]. We use a Monsoon meter to measure the current supply of the whole LoRaWAN shield after properly jumping the power wires. Monsoon is a power meter that can deliver the real-time power consumption readings to a computer via a USB cable [43]. Table 2 shows the measurement results under different SFs. The results are close to RN2483's datasheet, showing that the shield's encapsulating and interfacing circuits consume little power.

A possible concern about LoRaWAN is its low data rate to Power Consumption Ratio (DPR), compared with other low-power radios. For instance, with SF7 in the EU868 MHz band, the DPR is 11 kbps/38.9 mA = 0.28 kbps/mA. In contrast, the DPR for Zigbee is 250 kbps/19.5 mA = 12.82 kbps/mA. However, the severity of this concern should be discriminated regarding the aimed communication range. This is illustrated by an example of moving x bits of data over a distance of L meters by multiple hops. The radio energy used to move the x bits over a hop is $(P_{Tx} + P_{Rx}) \cdot \frac{x}{v}$, where P_{Tx} and P_{Rx} are the transmitting and receiving powers, respectively; v is the link data rate in bps. Thus, the total energy used by the network's radios to move x bits over L meters is $(P_{Tx} + P_{Rx}) \cdot \frac{x}{v} \cdot \frac{L}{d}$, where d is the typical one-hop transmission range. Considering $L = 1$ km, we set $\frac{L}{d}$ to be 1 and 10 for LoRaWAN and Zigbee, respectively. Moreover, we set the data rate v to be 11 kbps and 250 kbps for LoRaWAN and Zigbee, respectively. After applying respective power consumption measurements, LoRaWAN's total radio energy consumption is 2.94 times of Zigbee's. Although the above simplistic energy consumption estimation does not consider other factors like nodes' processor energy consumption and MAC, the result underlines our understanding. While LoRaWAN consumes more energy than Zigbee, it substantially simplifies the control-plane network design due to its one-hop nature. Moreover, the concern of LoRaWAN's higher energy consumption can be mitigated by the fact that the control plane's traffic volume is much lower than the data plane's. For instance, as measured in §6, the number of CTP-SCDP's control-plane frames is just about 5% of its data-plane packets. Thus, we believe that, for the control-plane networks, the energy saving by using high-DPR but short-range radios is not worth sacrificing network simplicity.

4.2.3 Latency profiling. Under TDMA, the LoRaWAN radio can sleep to save energy while waiting for the next time slot. The time delays in awaking the radio and transmitting a frame are critical to the radio's sleep scheduling and clock synchronization required by TDMA, respectively. The latency in awaking the radio from the RPi using the shield's C++ API is measured. Fig. 9 shows the distribution of the awaking latency over 500 tests that give satisfactory statistical significance of

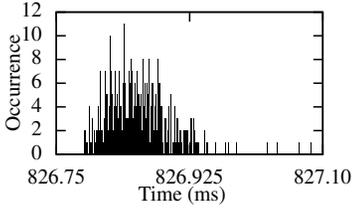


Fig. 9. Radio awaking latency.

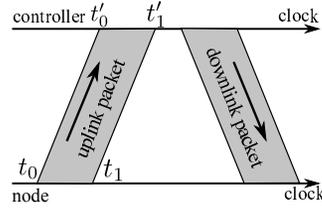


Fig. 10. A communication session.

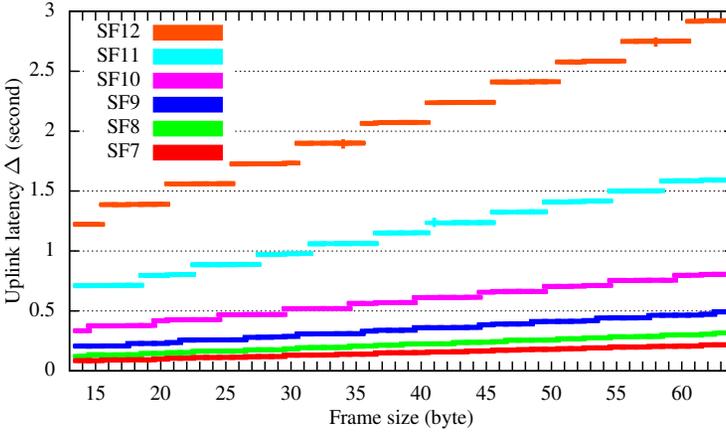


Fig. 11. Uplink latency under different SFs and frame sizes.

the measurement. The mean and standard deviation are 826.9 ms and 0.044 ms, respectively. The small standard deviation suggests that a LoRaCP node can awake the radio punctually for the next TDMA time slot.

Then, the latency in transmitting an uplink frame is measured. Fig. 10 illustrates the uplink transmission's timing. The node starts and completes the transmission when its clock values are t_0 and t_1 , respectively. The controller starts and completes the reception when its clock values are t'_0 and t'_1 , respectively. The t_0 , t_1 , and t'_1 can be recorded in the LoRaWAN shield's and concentrator's C++ user programs running at their RPis. To measure the uplink latency, the clocks of the node's and controller's RPis are synchronized using the Network Time Protocol (NTP) over an Ethernet that gives sub-ms synchronization accuracy. We define the uplink latency as $\Delta = t'_1 - t_0$.² Thus, the latency is determined by the data rate, which further depends on SF, and the frame size. Fig. 11 shows the box plots of the measured uplink latency under different SFs and frame sizes. As the latency has little variations under each setting, the boxes and whiskers of the plots are not visible. It can be seen that the latency increases with both frame size and SF, which are consistent with our understanding. Interestingly, for a certain SF, the latency exhibits step changes when the frame size increases. This is because each LoRa frame is a certain number of bits aligned for easy hardware handling. The above measurement results lay a foundation for developing LoRaWAN clock synchronization in §5.3.2.

²We do not use t_1 because it contains a non-negligible uncertain delay from the actual completion of the transmission to the LoRaWAN shield's C++ library's periodic pull of the event from the shield's hardware interface.

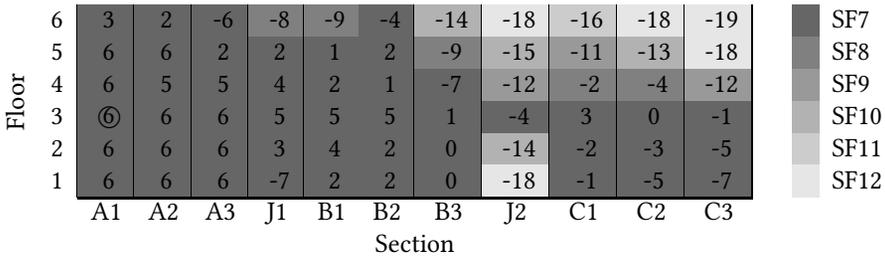


Fig. 12. SNR measurements in a 6-story building and the coverage of LoRaWAN with different SF settings. The numbers in the cells are the SNRs in dB measured by the RN2483 of the LoRaCP node. The LoRaCP controller is located in the cell indicated by a circle. The dimension over the sections is about 190 meters. With SF12, the whole building can be covered.

4.2.4 Indoor communication profiling. Various existing studies [35, 40] have investigated LoRaWAN's communication performance in outdoor environments. However, many MH-WSNs are deployed in indoor environments, e.g., a WirelessHART network for a manufacturing system and a building environment monitoring WSN. Thus, we profile the communication performance of LoRaWAN in a large 6-story concrete building. Along the building's long dimension of about 190 meters, the building has three sections (A, B, C) and two section junctions (J1 and J2). The LoRaCP controller is deployed in Section A1 on the third floor, as illustrated by the circle in Fig. 12. Then, a LoRaCP node is carried to different positions inside the building to measure the SNRs. Note that the Microchip RN2483 chip can estimate the SNR of the received signal. In each section, we measure three positions. The numbers in the cells of Fig. 12 are the SNR measurements in dB. It can be seen that the SNR measurements are from -19 dB to 6 dB. The SNR in general decreases with the distance between the LoRaCP node and the LoRaCP controller. Considerable SNR reductions are observed in the section junctions J1 and J2. This is because there are complex and steel structures in J1 and J2, such as stair cases, lifts, water facilities, and etc. For RN2483, the minimum SNRs required for reliable demodulation with SF7 to SF12 are -7.5 dB, -10 dB, -12.5 dB, -15 dB, -17.5 dB, and -20 dB, respectively [19]. Thus, in Fig. 12, a certain gray color is used to indicate the minimum SF setting that can cover each cell. It can be seen that most cells can be covered by a one-hop LoRaWAN with SF7. The whole building can be covered by a one-hop LoRaWAN with SF12. The above results show that although the indoor structures (e.g., concrete floors, walls, and steel objects) cause significant signal attenuation, with a large SF setting, a one-hop LoRaWAN network can well cover a building. For a larger building, multiple LoRaWAN gateways can be deployed to make sure every LoRaCP node in the building can communicate with the LoRaCP controller over a single wireless hop only.

4.2.5 Discussion. From the measurement results in [35], with line of sight in outdoor environments, LoRaWAN can achieve nearly 100% frame reception ratio when the distances are within 5 km. However, with obstruction, the frame reception will be severely affected. Thus, to apply the proposed one-hop out-of-band control plane design for outdoor MH-WSNs, extra care is needed to make sure each LoRaCP node can communicate with the LoRaCP controller reliably. Before system deployment, a remote sensing approach proposed in [15] can be used to predict the link quality and plan the LoRaCP node and LoRaWAN gateway positions. After system deployment, more LoRaWAN gateways can be added to fix coverage issues.

4.3 Performance Profiling of TI Sub-GHz Radio

In addition to LoRaWAN, this section investigates the feasibility of building one-hop out-of-band control planes using TI's SimpleLink sub-GHz long-range radio.

4.3.1 Introduction and characteristics of SimpleLink. SimpleLink [27] is a wireless MCU family made by Texas Instruments (TI). The SimpleLink covers a spectrum of wireless technologies including Wi-Fi, BLE, Thread, Zigbee, and sub-GHz radios. A SimpleLink platform also integrates an MCU to run user-defined programs. While TI provides a unified proprietary software development kit, the SimpleLink platforms can also run open-source operating systems such as Contiki-NG [13] and RIOT-OS [52]. Among various SimpleLink chips, the CC1350 and CC1352R are best suitable for building the one-hop out-of-band control planes. The CC1350, which has been adopted by the SensorTag platform, consists of a Cortex-M3 MCU, a BLE radio, and a sub-GHz long-range radio. The more recent CC1352R consists of a Cortex-M4F MCU, a sub-GHz long-range radio, and a 2.4GHz radio that supports multiple protocols including Zigbee and BLE. CC1352R has two operating modes for its sub-GHz radio: 2-GFSK mode and SimpleLink long-range mode. The 2-GFSK mode follows the IEEE 802.15.4g standard and adopts a binary Gaussian frequency shift keying modulation that can achieve 50 kbps and 200 kbps data rates with low and high output powers. The SimpleLink long-range mode uses the Direct-Sequence Spread Spectrum (DSSS) modulation that can achieve a maximum speed of 20 kbps (which is similar to LoRa) and a maximum communication range of 20 km. The form factor of the CC1352R is just $7\text{ mm} \times 7\text{ mm} \times 0.5\text{ mm}$. This shows that a miniature platform such as CC1352R may meet the hardware requirements of our proposed one-hop out-of-band control plane design. Since the CC1352R supports Zigbee, it can leverage on the legacy programs of TinyOS and Contiki that are designed for Zigbee, e.g., CTP. Thus, the CC1352R is a promising platform for implementing our proposed one-hop out-of-band control plane design.

In what follows, the energy consumption and latency of CC1352R are profiled using two pre-production launchpads from TI. Fig. 8 shows a CC1352R launchpad.

4.3.2 Power profiling. From the datasheet, CC1352R's current draw is 24.3 mA at 3.7 V when using the maximum transmitting power (14 dBm). Thus, the transmission power consumption is $24.3 \times 3.7 = 89.91\text{ mW}$. In the profiling, a voltage of 3.3 V is supplied to the CC1352R launchpad. We use the Monsoon power meter to measure the current draw of the launchpad. Due to the tight integration of the CC1352R chip and the launchpad, we cannot measure the sole power consumption of the CC1352R chip. Monsoon reads 125.8 mA and 161.7 mA when the launchpad is idle and transmitting using the SimpleLink mode, respectively. Thus, our estimate of CC1352R's power consumption in transmission is $(161.7\text{ mA} - 125.8\text{ mA}) \times 3.3\text{ V} = 118.47\text{ mW}$, which is about 20 mW higher than the datasheet value. A potential reason is that, when the CC1352R chip is in transmission, some supporting circuits on the launchpad are also active and consuming power. Using the same approach, the CC1352R's power consumption in the receiving mode is estimated as $11.1\text{ mA} \times 3.3\text{ V} = 36.63\text{ mW}$. In our future work, the power profiling approaches adopted in [8, 9] can be used to understand CC1352R launchpad's power consumption better.

4.3.3 Timing performance profiling. This section profiles the timing performance of CC1352R's sub-GHz radio. Different from the method for measuring the latency of LoRaWAN radio in transmitting a frame, we measure the error in synchronizing the internal clocks of two CC1352R nodes using their sub-GHz radios. Specifically, the two CC1352R nodes are placed close to each other and let them perform a round-trip communication as in NTP. The standard deviation of the one-way transmission delay (i.e., half of the round-trip time) characterizes the error in synchronizing the two nodes. We implement the round-trip timing in TI's Real-Time Operating System (TI-RTOS).

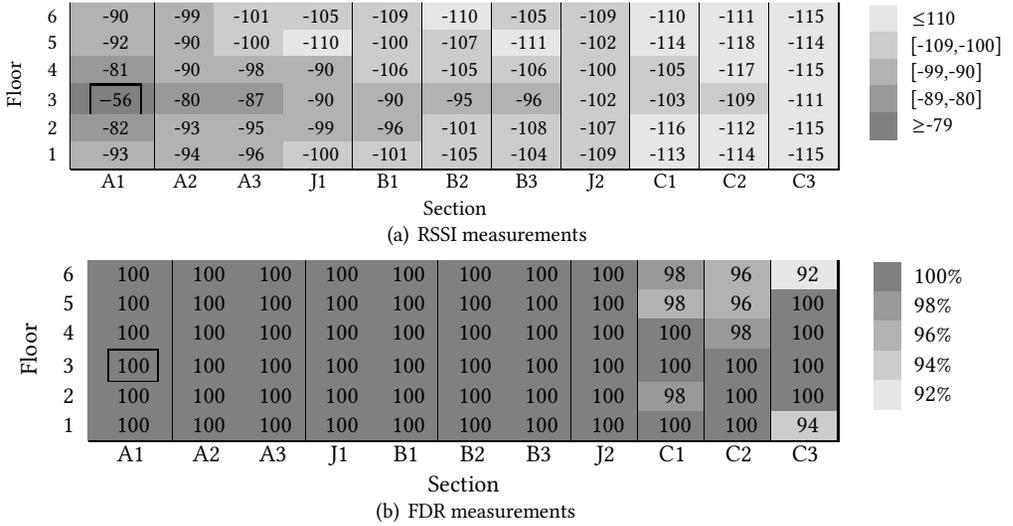


Fig. 13. RSSI and FDR measurements in a 6-story building. The number in a cell is the RSSI in dBm or FDR in percentage measured by a CC1352R in the cell. Another CC1352R is located in the cell indicated by a box. The dimension over the sections is about 190 meters.

In TI-RTOS, the transmission and arrival of the frames are timestamped using Radio Timer (RAT), which is a functional unit on the radio core providing high-resolution timing. From our measurements, the one-way transmission delay has a mean of $532.59 \mu\text{s}$ and a standard deviation of $2.50 \mu\text{s}$. Thus, the clock synchronization using CC1352R's sub-GHz radio can achieve microseconds accuracy. Note that as measured in §4.2.3, LoRaWAN's frame transmission timing has millisecond-level uncertainty, which results in the milliseconds clock synchronization error as shown in §5.3.2. The higher synchronization accuracy achieved by CC1352R is due to its hardware-level timestamping capability provided by RAT. In contrast, the LoRaWAN frame timestamping performed on the RPi is subjected to the overhead of the Raspbian operating system. These results show that CC1352R has good performance in timing the transmissions of frames, which is desirable for implementing the TMDA-based control planes.

4.3.4 Indoor communication profiling. To profile the indoor communication performance, a CC1352R node is fixed at a location and another CC1352R node is moved in the 6-story building mentioned in §4.2.4. Fig. 13 shows the Received Signal Strength Indicator (RSSI) and Frame Delivery Ratio (FDR) measurements when the mobile CC1352R node is in different locations in the building. In this set of measurements, both nodes adopt a transmission power of 14 dBm, same as the setting used for LoRaWAN. It can be seen from the results that the RSSI attenuates with distance. The CC1352R can achieve satisfactory FDRs throughout the building. These results show that a one-hop SimpleLink network can well cover a building.

4.3.5 Discussion. From the above results, the power, latency, and indoor communication performance profiles of CC1352R's SimpleLink long-range radio are similar to LoRaWAN's. Moreover, CC1352R has a built-in 2.4GHz radio that supports Zigbee. Thus, CC1352R is a promising platform for implementing the proposed one-hop out-of-band control plane design. As CC1352R is still in the pre-production phase, this paper focuses on using the LoRaCP hardware prototype presented

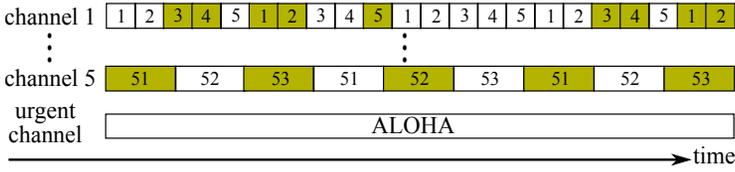


Fig. 14. Illustration of LoRaCP-TxC (shaded blocks mean heartbeat slots).

in §4.2.1 to implement and evaluate our proposed approach. The implementation of the out-of-band control planes using CC1352R will be similar.

5 DESIGN AND IMPLEMENTATION OF LoRaCP

5.1 System Overview and LoRaCP-TxC

The goal of LoRaCP is to use LoRaWAN's uplinks and downlinks to transmit *network reports* from the nodes to the controller and *network commands* from the controller to the nodes, respectively. The network commands have two categories: a *reactive network command* to a node is in response to a precedent network report from the node, whereas an *active network command* is initiated by the controller. All the control-plane transmissions are managed by LoRaCP's transmission control protocol as illustrated in Fig. 14, which is called *LoRaCP-TxC*. As discussed in §4.1, LoRaWAN has six concurrent uplink channels. Five of them use TDMA, while the remaining one (called *urgent channel*) uses ALOHA to transmit urgent frames. The five concurrent TDMA channels increase the throughput for the network reports. The urgent channel mitigates the rigidity of TDMA and allows the control-plane application developers to deal with urgent situations such as sudden strong interference or even malicious jamming to the data-plane network. As the TDMA channels have different data rates, their time slot lengths can be different to achieve the same maximum frame size. The time slots of a TDMA channel are allocated in a round-robin fashion to the LoRaCP nodes that use the channel. The LoRaCP nodes can be assigned to the TDMA channels to balance their time delays in waiting for the next time slot, while considering the channels' communication ranges and the nodes' distances to the controller.

Certain regions impose duty cycle requirements on the sub-GHz ISM bands used by LoRaWAN. For instance, in Europe, a LoRaWAN end device operating in the EU868 MHz band needs to conform to a duty cycle requirement specified by The European Telecommunications Standards Institute [28]. The duty cycle upper limit can be 0.1%, 1% or 10% depending on the used sub-band. Other regions impose dwell time, i.e., the channel occupancy time. For instance, in North America, the dwell time of the LoRaWAN operating on a 125 kHz bandwidth centered at any frequency must not be longer than 0.4 seconds within any 20-second period [11]. Note that a dwell time requirement can be translated to a duty cycle requirement. For instance, the dwell time requirement mentioned earlier can be translated to a 2% duty cycle requirement. To meet the duty cycle requirement, each node has a minimum waiting time between two consecutive transmissions. Accordingly, the time slot lengths of the channels in LoRaCP-TxC can be designed to meet the duty cycle requirement. Specifically, the time slot length for a channel should satisfy

$$\text{time slot length} \geq \max \left\{ \frac{\text{a node's minimum waiting time}}{\text{the number of nodes assigned to the channel}}, 3 \text{ seconds} \right\},$$

where the 3-second time is the minimum time for completing a communication session consisting of an uplink and two optional downlink transmissions. Note that a node will open two continuous

1-second windows waiting for the downlink transmission. The remaining one second in a communication session is for the uplink transmission and data processing. The minimum waiting time can be derived based on the LoRa's on air time calculation method described in [53]. For instance, for the channel using SF7, bandwidth of 125 kHz and coding rate of 4/5, if the payload size is 30 bytes and the duty cycle requirement is 1%, the minimum waiting time is 8.73 seconds. If there are 5 LoRaCP nodes in this channel, the time slot length can be set as $\max\left\{\frac{8.73}{5}, 3\right\} = 3$ seconds.

We now present two features of LoRaCP-TxC that address LoRaWAN's downlink-uplink asymmetry and lossy links.

- **Heartbeat time slots:** When a node has no uplink data to transmit, it can skip its next time slot. However, because any downlink frame must be in response to a precedent uplink frame, LoRaCP-TxC designates periodic heartbeat time slots for each node. For instance, in Fig. 14, the shaded blocks represent heartbeat slots. In Channel 1, the heartbeat period is three time slots. A node must transmit an uplink frame in a heartbeat slot. This opens a downlink window to maintain the clock synchronization of the node (cf. §5.3.2) and send active network commands. The heartbeat period can be set according to the nodes' clock drift rates and the required clock accuracy to avoid TDMA panic. The heartbeats also help the LoRaCP controller be aware of whether a node is still alive.
- **Negative acknowledgment (NAK):** To deal with frame losses, acknowledging all concurrent uplink transmissions is wasteful because of the downlink-uplink asymmetry. Thus, LoRaCP uses the NAK scheme. In LoRaWAN, the uplink and downlink frames from/to a node have continuously increasing counters, respectively. Thus, both the controller and the nodes can detect if there are lost frames by checking the continuity of the frame counters. If the controller detects lost frames, it sends an NAK using the subsequent downlink transmission to notify the node, which can then use the urgent channel or wait for the next TDMA slot to resend the lost data. The node can also send NAK using the urgent channel or the next TDMA slot to request lost frames. With the NAK scheme, the controller does not need to respond to a node's network report if there are no network commands for the node and no lost frames. This design mitigates the contention for the downlink time.

5.2 Software Architectures of LoRaCP Node and Controller

In §4.2, we have introduced the hardware prototypes of the LoRaCP node and controller. This section presents their software architectures as illustrated in Fig. 15. Note that the software of LoRaCP, including the transmission control protocol LoRaCP-TxC, is implemented in the application layer using the program library of the used LoRaWAN platform. The LoRaWAN's ALOHA MAC is beneath the program library.

5.2.1 LoRaCP node. A C++ forwarder program *LoRaCPFwd* runs on the RPi to buffer and forward the data between Kmote and the LoRaWAN shield, while following LoRaCP-TxC. The node parts of the clock synchronization and TDMA are also implemented in *LoRaCPFwd*. The Kmote runs TinyOS. We design a TinyOS module *LoRaCPC* that provides the *AMSend* and *Receive* interfaces to send and receive data to/from the RPi through serial communications. Thus, in our prototype design, the Kmote uses LoRaWAN as a service.

5.2.2 LoRaCP controller. The RPi of the controller runs an open-source LoRaWAN server architecture [10] consisting of *packet_forwarder*, *LoRa Gateway Bridge*, *LoRa Server*, and *LoRa App Server*. This architecture, through providing JSON-based APIs to subscribe/send messages from/to the LoRaWAN network, greatly simplifies the design of centralized network control applications. The role of this architecture is similar to that of an SDN controller platform (e.g., OpenDaylight)

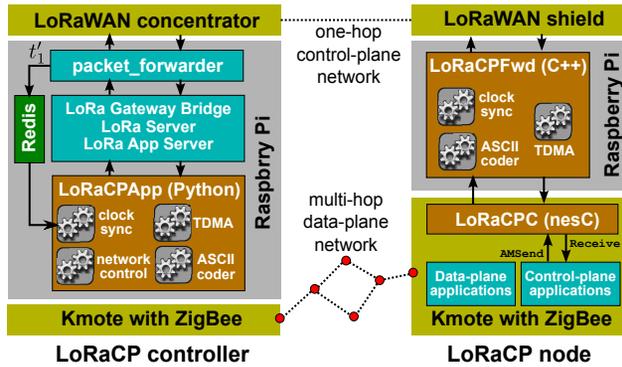


Fig. 15. Software architectures of LoRaCP controller and node. (The illustration includes a Zigbee radio for the controller to be a control-plane sink.)

that facilitates the design of SDN control applications. In this paper, the centralized network controls and the controller parts of the clock synchronization and TDMA are implemented in a single Python program called *LoRaCPApp*. Note that the LoRaWAN server architecture [10] supports multiple LoRaWAN gateways. Although this paper focuses on a single LoRaCP controller, in the future work, the multi-gateway support can be exploited to develop redundant LoRaCP controllers to improve the system's reliability against a single point of failure.

5.3 Implementation of LoRaCP Components

This section provides implementation details of LoRaCP's integer coding scheme, clock synchronization, and TDMA.

5.3.1 ASCII coding of integers. ASCII code words are used to directly encode integers for better efficiency. The use of the ASCII code words is because LoRa admits hexadecimal ASCII string only. Since there are 128 ASCII code words, the last one is reserved for separator and the remaining 127 code words to encode an integer. Specifically, the decimal integer is presented in the base-127 numeral system. For instance, $10000 = 78 \times 127^1 + 94 \times 127^0 = (78, 94)_{127}$. Then, the hexadecimal ASCII string consisting of the digits of the base-127 representation is fed to LoRa. For the above example, the string is "4E5E", much shorter than the "3130303030" in LoRa's default coding. The decoding is simply the reverse process. The above ASCII code is implemented in *LoRaCPFwd* and *LoRaCPApp* at the LoRaCP node and controller, respectively.

5.3.2 Clock synchronization. Clock synchronization is a basis for implementing TDMA. Although there are various existing clock synchronization protocols for MH-WSNs (e.g., FTSP), if the LoRaCP nodes are synchronized to the controller using the data-plane network, the control plane's TDMA will depend on the data-plane network, incurring the undesirable coupling. Thus, the LoRaCP nodes should be synchronized to the controller using the control-plane network. However, there is still limited research on clock synchronization over LoRaWAN. In our prototype system, the RPi's clock is used as the node's or controller's clock. Although the LoRaWAN devices and the K mote have their own timers, using the RPi's clock can simplify the evaluation of the accuracy of the LoRaWAN clock synchronization using the RPi's Ethernet interface.

To save the downlink time, LoRaCP does not prescribe dedicated frames for clock synchronization. Instead, LoRaCP piggybacks several bytes to each control-plane frame for clock synchronization. Specifically, each uplink frame is appended with the node's clock value t_0 as illustrated in

Fig. 10. The controller records its clock value t'_1 on completion of the frame reception. The clock offset between the node and the controller, denoted by δ , can be estimated as $\delta = t'_1 - (t_0 + \Delta)$, where Δ is the uplink latency presented in Fig. 11. Then, the controller piggybacks δ onto the downlink frame as illustrated in Fig. 10. Upon receiving δ , the node resets its clock by $t = t + \delta$, where t denotes the node's current clock value. Alternatively, the node's clock advance speed can be calibrated according to δ using a negative feedback loop.

We now discuss several implementation issues of the above clock synchronization approach. First, the LoRaWAN frame header added by the shield has changeable size because the integers in the headers are represented as variable length hexadecimal ASCII strings (cf. §4.1.2). As shown in Fig. 11, the uplink latency Δ has a complex relationship with the frame size in different channels. When the LoRaCP controller receives the uplink frame, it checks the actual frame size and the SF used by the node to query the corresponding Δ from the data in Fig. 11. Thus, for LoRaWAN clock synchronization, the prior knowledge in Fig. 11 is critical. Note that most MH-WSN clock synchronization approaches are free from this frame size dependence issue because they use dedicated synchronization frames with fixed sizes or the frame size has little impact on transmission latency. Second, we modify *packet_forwarder*, i.e., LoRaWAN concentrators' driver program, to record t'_1 , because other components of the LoRaWAN server architecture may suffer software delays. As illustrated in Fig. 15, the timestamp t'_1 , together with the corresponding source ID and frame ID, are written into a Redis in-memory database and then retrieved by the *LoRaCPApp* to compute δ .

The synchronization accuracy of the above approach is measured using the `ntupdate` tool to check the clock offset between the node and the controller over a local Ethernet network connecting the RPIs. The mean absolute synchronization error is 2.9 ms with a standard deviation of 1.7 ms. Given the second-level frame transmission time, such synchronization errors of a few milliseconds are satisfactory.

5.3.3 TDMA. The prototype LoRaCP node controls the sleep of the LoRaWAN radio and transmissions of frames based on its RPI's synchronized clock. Specifically, if *LoRaCPFwd* has received a network report from the Kmote, the RPI starts awaking the LoRaWAN radio 850 ms before its next TDMA time slot, transmits the report in the time slot, receives any subsequent network command, re-transmits frames using the urgent channel if a NAK is received. Finally, *LoRaCPFwd* forwards all received network commands to the Kmote. In our current experimental implementation, LoRaWAN channels and time slots are assigned to nodes manually.

6 PERFORMANCE EVALUATION

Various testbed experiments are conducted to evaluate our LoRaCP implementation.

6.1 Experiment Methodology and Settings

LoRaCP is applied to implement the out-of-band CTP-SCDP presented in §3. Specifically, if the Kmote of a LoRaCP node detects a change of ETX with any of its neighbor node, it uses the *LoRaCPC* to send the latest ETX using a network report frame to the LoRaCP controller. Upon receiving the ETX update, the controller's *LoRaCPApp* python program computes the optimal routing and pushes network commands containing new parent node information to the downlink queue of the LoRaWAN server architecture. Upon receiving a network command, a LoRaCP node updates its parent node accordingly. In the data plane, each node generates a data packet every eight seconds.

We conduct experiments on a testbed consisting of a LoRaCP controller and 15 LoRaCP nodes. The LoRa modules use the 868 MHz ISM band. The nodes are placed at the grid points of a lab space.

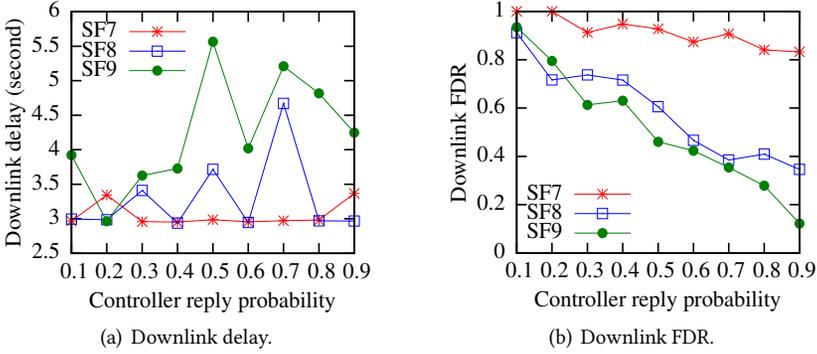


Fig. 16. Control plane communication performance.

The nodes are evenly divided to use three LoRaWAN channels (SF7, SF8, and SF9). The time slot lengths in these three channels are 3, 4, and 5 seconds, respectively.³ The controller uses the first downlink window RX1 to transmit network commands. Before the RX1 window, the controller has a *wait time* of one second to compute the network commands, which is generally sufficient. On our 16-node testbed, each LoRaCP has a time slot every 25 seconds or less. For larger networks, to maintain this rotating period for each node, multiple geographically distributed nodes in the same channel can be assigned to use the same time slot, since they unlikely report ETX changes at the same time.

6.2 Experiment Results

We conduct three sets of experiments: §6.2.1 evaluates the control plane communication performance; §6.2.2 evaluates the control plane performance of CTP-SCDP; §6.2.3 compares CTP and CTP-SCDP.

6.2.1 Control plane communication performance. While the concurrent uplink channels increase the throughput for network reports, LoRaWAN's downlink-uplink asymmetry presents a bottleneck for the downlink communications. The downlink performance is evaluated. Specifically, each LoRaCP node transmits a network report every its time slot. Thus, the controller receives frames from the three channels concurrently almost at all the time. It replies to each network report with a certain probability. The frame size of the replies ranges from 29 to 33 bytes. NAK is turned off in these tests.

Fig. 16 shows the average control-plane downlink delays and Frame Delivery Ratios (FDRs) of different channels versus the probability that the controller replies. The downlink delay is measured as the time duration between i) the controller's *LoRaCPApp* pushes a network command to the LoRaWAN server architecture and ii) the node's *LoRaCPFwd* receives the command. This downlink delay includes the wait time of one second. From Fig. 16(a), the average downlink delay does not significantly increase with the controller's reply probability. The average delay ranges from 3 seconds to 5.5 seconds. It increases with the SF, because a larger SF has a lower data rate. Fig. 16(b) shows the control-plane downlink FDR versus the controller's reply probability. The FDR

³The duty cycles in the three channels are 0.9%, 1.4%, 2.1%. Note that the region in which this paper's experiments are conducted does not impose duty cycle requirement on the 868 MHz ISM band [12]; it only imposes a transmitting power upper limit. To meet Europe's 1% duty cycle requirement, the approach presented in §5.1 can be used to configure the time slot lengths.

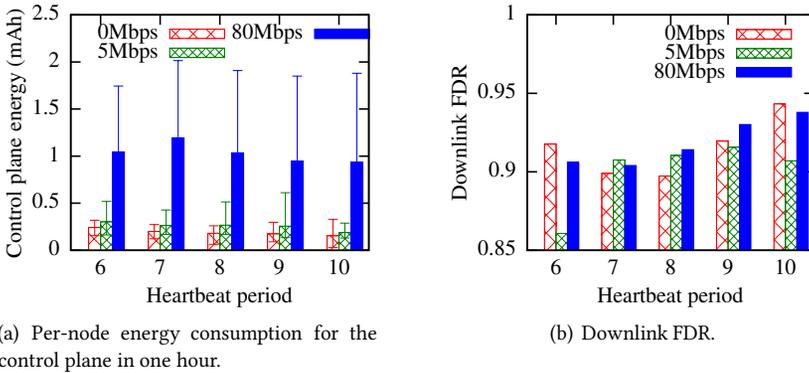


Fig. 17. CTP-SCDP control plane performance under Wi-Fi interference against the data-plane network. The error bar represents min and max values.

decreases with the reply probability. This is because the open-source LoRaWAN server architecture [10] drops frames when it receives excessive frames to be transmitted beyond the downlink throughput. From the results in Fig. 16, the downlink bottleneck mainly affects the downlink FDR. Thus, in the remaining experiments, the downlink FDR is used to assess whether the control plane performance is throttled by the downlink-uplink asymmetry.

6.2.2 Control plane performance in CTP-SCDP. The performance of CTP-SCDP's control-plane network is evaluated. To create data-plane link quality variations, a laptop placed close to the testbed is used to generate Wi-Fi traffic to interfere with the Zigbee data-plane network. Zigbee radios use Channel 18 and the Wi-Fi AP uses Channel 6, which interfere with each other. On the laptop, `iperf3` is used to generate data traffic at a specified bit rate. This experiment methodology well captures the increasingly crowded 2.4 GHz ISM band used by the Zigbee-/BLE-based data-plane networks. In the presence of the Wi-Fi interference, the CTP-SCDP generates more control-plane messages to report the volatile link ETXes of the data-plane network to the LoRaCP controller.

First, we estimate the energy consumption of each LoRaCP node's LoRaWAN shield by multiplying the transmitting/receiving currents with the measured total times in respective modes. Fig. 17(a) shows the error bars of per-node energy consumption by the shield in one hour under different settings of heartbeat period and Wi-Fi interference intensity. The control-plane energy consumption increases with the interference intensity due to the increased control-plane messages. When we do not generate Wi-Fi interference, the energy consumption decreases with the heartbeat period. This is because, in the absence of the interference, the link ETXes seldom change and most control-plane messages are the heartbeats. In the presence of interference (i.e., 5 Mbps and 80 Mbps), the energy consumption has no monotonic relationship with the heartbeat period, because the node will utilize the non-heartbeat time slots to report the volatile ETXes. From Fig. 17(a), with no and intensive interference (80 Mbps), the per-node power consumption by the control plane averaged over time is about 0.25 mA and 1 mA, respectively, which are comparable to or lower than the power consumption of low-power microcontrollers (MCUs). For instance, the active power of TelosB's MCU is 1.8 mA, whereas the recent Firestorm's MCU consumes 8.6 mA in the common configuration [5].

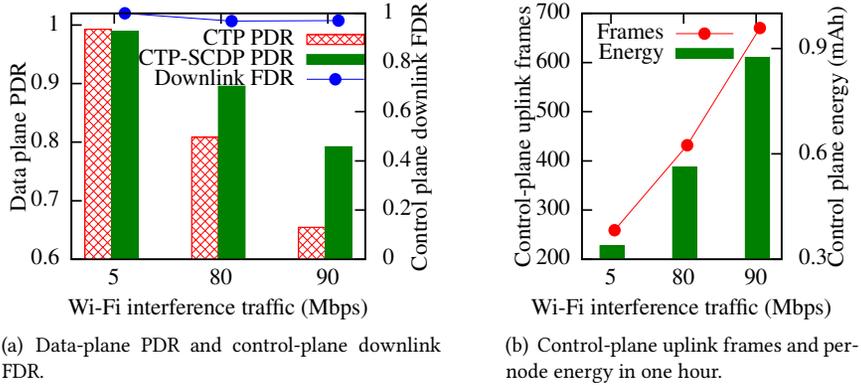


Fig. 18. Performance comparison between CTP and CTP-SCDP.

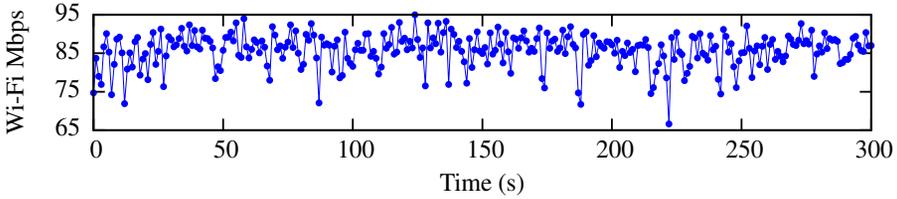


Fig. 19. Wi-Fi data rate fluctuations over time (setpoint: 90 Mbps).

Second, we measure the average control-plane downlink FDR over all channels. The results are shown in Fig. 17(b). Even if the data-plane network experiences intensive interference, the FDR is generally above 90%. Thus, the CTP-SCDP's control plane is still beyond the downlink bottleneck.

6.2.3 Comparison between CTP and CTP-SCDP. We load CTP to eight nodes and CTP-SCDP to another eight nodes. CTP and CTP-SCDP run side by side on the testbed, so that they experience almost the same Wi-Fi interference for fair comparison. CTP-SCDP's LoRaCP heartbeat period is 10. Fig. 18(a) shows the data plane's PDR, i.e., the ratio of the Zigbee packets received by the data-plane sink over all packets generated by the source nodes. When the Wi-Fi interference intensity is low (5 Mbps), CTP and CTP-SCDP achieve similarly high PDRs. When the interference intensity is high (80 Mbps), CTP-SCDP's PDR is 10% higher than CTP's. When the interference intensity is 90 Mbps, CTP's PDR drops to 65%, while CTP-SCDP's is 80%. Note that the actual data rate of the Wi-Fi interference traffic fluctuates over time. Fig. 19 shows the actual interference data rate when the setpoint to `iperf3` is 90 Mbps. Moreover, the fluctuation level increases with the setpoint. The data rate deviations are 0.8 Mbps only and up to 20 Mbps for setpoints 5 Mbps and 90 Mbps, respectively. Thus, the control-plane networks experience more dynamic interference with a larger setpoint, resulting in the increasing PDR gain of CTP-SCDP over CTP with the interference intensity setting. This result is consistent with our observation from the simulation study in §3.1 that CTP cannot handle dynamic network conditions well.

Fig. 18(a) also shows the control-plane downlink FDRs, which are above 97%. This suggests that the control plane is beyond the downlink bottleneck. Fig. 18(b) shows the total number of control-plane uplink frames of CTP-SCDP during one hour and the projected per-node energy consumption by the LoRaWAN shield. In the presence of stronger interference, more uplink frames

will be transmitted to report the volatile ETXes. With 5 Mbps and 90 Mbps interference, the total numbers of data-plane transmissions (including beacons and forwarded packets) are 5,022 and 10,024, respectively. The corresponding numbers of control-plane uplink frames are just 5.2% and 6.7% of these data-plane transmissions. With strong interference (90 Mbps), the per-node control-plane power consumption averaged over time is less than 0.9 mA, consistent with the results in Fig. 17 obtained with 15 nodes.

7 CONCLUSION AND FUTURE WORK

This paper studied how to use Low-Power Wide-Area Network (LPWAN) radios to form one-hop out-of-band control planes for multi-hop wireless sensor networks through extensive simulational and hardware profiling studies, system design, and testbed evaluation. Through an example application of Collection Tree Protocol (CTP), the simulational study showed the advantages of the one-hop out-of-band control plane design over the distributed network control scheme and the in-band control plane design. Our hardware profiling study based on two LPWAN technologies, i.e., LoRaWAN and CC1352R's sub-GHz long-range radio, showed their characteristics that should be considered in the design of the out-of-band control planes. The designed LoRaWAN-based control plane, LoRaCP, is applied to physically separate the control plane of CTP from its Zigbee-based data-plane network. Experiments showed that LoRaCP increases CTP's packet delivery ratio from 65% to 80% in the presence of external interference, while consuming a per-node average radio current of 0.9 mA only with an operating voltage of 3.3 V. This current consumption is comparable to or lower than those of low-power MCUs found on sensor network platforms such as TelosB and Firestorm.

In the future work, it is interesting to fully implement one-hop out-of-band control planes based on TI's sub-GHz radios. As the Contiki-NG operating system has better support of new hardware, the design of a Contiki-NG code library that provides well defined interfaces to facilitate the implementation of one-hop out-of-band control planes will be of great interest and meaningful contribution.

ACKNOWLEDGMENTS

The authors acknowledge Dusit Niyato for constructive comments during the early stage of this research.

REFERENCES

- [1] 2015. *The Smart Economy: The Internet of Everything*. Technical Report. Semico Research & Consulting Group.
- [2] 2018. TinyOS Production. Retrieved Jun 3, 2019 from <https://github.com/tinyprod>
- [3] 2019. iC880A - LoRaWAN Concentrator 868MHz. Retrieved Jun 3, 2019 from <https://wireless-solutions.de/products/radiomodules/ic880a.html>
- [4] Atul Adya, Paramvir Bahl, Jitendra Padhye, Alec Wolman, and Lidong Zhou. 2004. A multi-radio unification protocol for IEEE 802.11 wireless networks. In *The 1st International Conference on Broadband Networks (BroadNets)*. IEEE, 344–354.
- [5] Michael P Andersen, Gabe Fierro, and David E Culler. 2016. System design for a synergistic, low power mote/ble embedded platform. In *The 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 1–12.
- [6] Paramvir Bahl, Atul Adya, Jitendra Padhye, and Alec Walman. 2004. Reconsidering wireless systems with multiple radios. *ACM SIGCOMM Computer Communication Review* 34, 5 (2004), 39–46.
- [7] Inc. Bluetooth SIG. 2019. Bluetooth LE: mesh. Retrieved Jun 3, 2019 from <https://www.bluetooth.com/bluetooth-technology/topology-options>
- [8] Pietro Boccadoro, Michele Barile, Giuseppe Piro, and Luigi Alfredo Grieco. 2016. Energy consumption analysis of TSCH-enabled platforms for the Industrial-IoT. In *IEEE 2nd International Forum on Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI)*. IEEE, 1–5.

- [9] Pietro Boccadoro, Giuseppe Piro, Domenico Striccoli, and Luigi Alfredo Grieco. 2018. Experimental Comparison of Industrial Internet of Things Protocol Stacks in Time Slotted Channel Hopping Scenarios. In *IEEE International Conference on Communications (ICC)*. IEEE, 1–6.
- [10] Orne Brocaar. 2019. LoRa server system architecture. Retrieved Jun 3, 2019 from <https://www.loraserver.io/overview/architecture/>
- [11] Federal Communications Commission. 2019. FCC Regulations for ISM Band Devices: 902 - 928 MHz. Retrieved Jun 3, 2019 from https://www.semtech.com/uploads/documents/fcc_part15_regulations_semtech.pdf
- [12] Info communications Media Development Authority of Singapore. 2017. Spectrum Management Handbook (Issue 1 Rev 2.9 - July 2017). Retrieved June 3, 2019 from <https://bit.ly/2QL6jr5>
- [13] Contiki. 2019. Contiki-NG. Retrieved Jun 3, 2019 from <http://www.contiki-ng.org/>
- [14] Peter Dely, Andreas Kessler, and Nico Bayer. 2011. OpenFlow for wireless mesh networks. In *The 20th International Conference on Computer Communications and Networks (ICCCN)*. IEEE, 1–6.
- [15] Silvia Demetri, Marco Zúñiga, Gian Pietro Picco, Fernando Kuipers, Lorenzo Bruzzone, and Thomas Telkamp. 2019. Automated Estimation of Link Quality for LoRa: A Remote Sensing Approach. In *The 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*.
- [16] Adwait Dongare, Revathy Narayanan, Akshay Gadre, Anh Luong, Artur Balanuta, Swarun Kumar, Bob Iannucci, and Anthony Rowe. 2018. Charm: exploiting geographical diversity through coherent combining in low-power wide-area networks. In *The 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 60–71.
- [17] Gholamhossein Ekbatanifard, Philipp Sommer, Branislav Kusy, Venkat Iyer, and Koen Langendoen. 2013. Fastforward: High-throughput dual-radio streaming. In *IEEE 10th International Conference on Mobile Ad-Hoc and Sensor Systems (MASS)*. IEEE, 209–213.
- [18] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yağan. 2017. Empowering Low-Power Wide Area Networks in Urban Settings. In *The Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*. ACM, 309–321.
- [19] Modtronix Engineering. 2018. Wireless SX1276 LoRa Module. Retrieved Jun 3, 2019 from <http://modtronix.com/inair9b.html>
- [20] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, David Moss, and Philip Levis. 2009. Collection tree protocol. In *The 7th ACM conference on embedded networked sensor systems (SenSys)*. ACM, 1–14.
- [21] Karim Habak, Khaled A Harras, and Moustafa Youssef. 2015. Bandwidth aggregation techniques in heterogeneous multi-homed devices: A survey. *Computer Networks* 92 (2015), 168–188.
- [22] Mehrdad Hesar, Ali Najafi, and Shyamnath Gollakota. 2019. NetScatter: Enabling Large-Scale Backscatter Networks. In *The 16th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. USENIX.
- [23] Huawei Huang, Song Guo, Weifa Liang, Keqiu Li, Baoliu Ye, and Weihua Zhuang. 2016. Near-optimal routing protection for in-band software-defined heterogeneous networks. *IEEE Journal on Selected Areas in Communications (J-SAC)* 34, 11 (2016), 2918–2934.
- [24] Texas Instruments Incorporated. 2018. CC1352R. Retrieved Jun 3, 2019 from <http://www.ti.com/product/CC1352R>
- [25] Texas Instruments Incorporated. 2018. TI CC1350. Retrieved Jun 3, 2019 from <http://www.ti.com/product/CC1350>
- [26] Texas Instruments Incorporated. 2018. TI SensorTag. Retrieved Jun 3, 2019 from http://www.ti.com/ww/en/wireless_connectivity/sensortag/
- [27] Texas Instruments Incorporated. 2019. SimpleLink. Retrieved Jun 3, 2019 from <http://www.ti.com/wireless-connectivity/simplelink-solutions/overview/overview.html>
- [28] The European Telecommunications Standards Institute. 2019. Technical characteristics for Low Power Wide Area Networks Chirp Spread Spectrum (LPWAN-CSS) operating in the UHF spectrum below 1 GHz. Retrieved Jun 3, 2019 from https://www.etsi.org/docdeliver/etsi_tr/103500_103599/103526/01.01.01_60/tr_103526v010101p.docx
- [29] Nachikethas A Jagadeesan and Bhaskar Krishnamachari. 2015. Software-defined networking paradigms in wireless networks: a survey. *ACM Computing Surveys (CSUR)* 47, 2 (2015), 27.
- [30] Anders R Jensen, Mads Lauridsen, Preben Mogensen, Troels B Sørensen, and Per Jensen. 2012. LTE UE power consumption model: For system level energy and performance optimization. In *2012 IEEE Vehicular Technology Conference (VTC Fall)*. IEEE, 1–5.
- [31] Srikanth Kandula, Kate Ching-Ju Lin, Tural Badirkhanli, and Dina Katabi. 2008. FatVAP: Aggregating AP Backhaul Capacity to Maximize Throughput. In *The 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, Vol. 8. 89–104.
- [32] Murad Kaplan, Chenyu Zheng, Matthew Monaco, Eric Keller, and Douglas Sicker. 2014. WASP: a software-defined communication layer for hybrid wireless networks. In *ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*. IEEE, 5–15.

- [33] Lorenzo Keller, Anh Le, Blerim Cici, Hulya Seferoglu, Christina Fragouli, and Athina Markopoulou. 2012. Microcast: Cooperative video streaming on smartphones. In *The 10th international conference on Mobile systems, applications, and services (MobiSys)*. ACM, 57–70.
- [34] Philip Levis, Nelson Lee, Matt Welsh, and David Culler. 2003. TOSSIM: Accurate and scalable simulation of entire TinyOS applications. In *The 1st international conference on Embedded networked sensor systems (SenSys)*. ACM, 126–137.
- [35] Jansen C. Liando, Amalinda Gamage, Agustinus W. Tengourtiou, and Mo Li. [n. d.]. Known and Unknown Facts of LoRa: Experiences from a Large Scale Measurement Study. *ACM Transactions on Sensor Networks (TOSN)* ([n. d.]). Accepted, in press.
- [36] Yeon-sup Lim, Yung-Chih Chen, Erich M Nahum, Don Towsley, Richard J Gibbens, and Emmanuel Cecchet. 2015. Design, implementation, and evaluation of energy-aware multi-path TCP. In *The 11th ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT)*. ACM, 30.
- [37] International Labmate Limited. 2012. WirelessHART Installed Networks Exceed 8,000 at Major Manufacturing Sites Worldwide. Retrieved Jun 3, 2019 from <https://bit.ly/2SiwSbu>
- [38] Tie Luo, Hwee-Pink Tan, and Tony QS Quek. 2012. Sensor OpenFlow: Enabling software-defined wireless sensor networks. *IEEE Communications Letters* 16, 11 (2012), 1896–1899.
- [39] Dimitrios Lymberopoulos, Nissanka B Priyantha, Michel Goraczko, and Feng Zhao. 2008. Towards energy efficient design of multi-radio platforms for wireless sensor networks. In *The 7th international conference on Information processing in sensor networks (IPSN)*. IEEE, 257–268.
- [40] Paul J Marcelis, Vijay S Rao, and R Venkatesha Prasad. 2017. DaRe: Data recovery through application layer coding for LoRaWAN. In *IEEE/ACM Second International Conference on Internet-of-Things Design and Implementation (IoTDI)*. IEEE, 97–108.
- [41] Microchip. 2019. RN2483. Retrieved Jun 3, 2019 from <https://www.microchip.com/wwwproducts/en/RN2483>
- [42] Microchip. 2019. RN2483: Low-Power Long Range LoRa Technology Transceiver Module. Retrieved Jun 3, 2019 from <https://bit.ly/2HOeuQF>
- [43] Inc. Monsoon Solutions. 2019. Monsoon Power Monitor". Retrieved Jun 3, 2019 from <http://msoon.github.io/powermonitor/PowerTool/doc/LVPM%20Manual.pdf>
- [44] Di Mu, Yunpeng Ge, Mo Sha, Steve Paul, Niranjana Ravichandra, and Souma Chowdhury. 2017. Adaptive radio and transmission power selection for Internet of Things. In *IEEE/ACM 25th International Symposium on Quality of Service (IWQoS)*. IEEE, 1–10.
- [45] Ashkan Nikravesh, Yihua Guo, Feng Qian, Z Morley Mao, and Subhabrata Sen. 2016. An in-depth understanding of multipath TCP on mobile devices: measurement and system design. In *The 22nd Annual International Conference on Mobile Computing and Networking (MobiCom)*. ACM, 189–201.
- [46] Shahriar Nirjon, Angela Nicoara, Cheng-Hsin Hsu, Jatinder Pal Singh, and John A Stankovic. 2014. MultiNets: A system for real-time switching between multiple network interfaces on mobile devices. *ACM Transactions on Embedded Computing Systems (TECS)* 13, 4s (2014), 121.
- [47] Bruno Astuto A Nunes, Marc Mendonca, Xuan-Nam Nguyen, Katia Obraczka, and Thierry Turletti. 2014. A survey of software-defined networking: Past, present, and future of programmable networks. *IEEE Communications Surveys & Tutorials* 16, 3 (2014), 1617–1634.
- [48] S.L. Openmote Technologies. 2018. Openmote. Retrieved Jun 3, 2019 from <http://www.openmote.com/>
- [49] Yao Peng, Longfei Shanguan, Yue Hu, Yujie Qian, Xianshang Lin, Xiaojiang Chen, Dingyi Fang, and Kyle Jamieson. 2018. PLoRa: a passive long-range data network from ambient LoRa transmissions. In *The 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*. ACM, 147–160.
- [50] T. PetriÄĀ, M. Goessens, L. Nuaymi, L. Toutain, and A. Pelov. 2016. Measurements, performance and analysis of LoRa FABIAN, a real-world implementation of LPWAN. In *IEEE 27th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*. 1–7.
- [51] Pycom. 2019. LoPy4. Retrieved Jun 3, 2019 from <https://pycom.io/product/lopy4/>
- [52] RIOT. 2013. RIOT: The friendly Operating System for the Internet of Things. Retrieved Jun 3, 2019 from <https://riot-os.org/>
- [53] Semtech. 2019. SX1272/3/6/7/8: LoRa Modem Designer’s Guide. Retrieved Jun 3, 2019 from <https://bit.ly/2WozGWO>
- [54] Semtech. 2019. SX1276/77/78/79 – 137MHz to 1020MHz Low Power Long Range Transceiver. Retrieved Jun 3, 2019 from https://www.semtech.com/uploads/documents/DS_SX1276-7-8-9_W_APP_V6.pdf
- [55] Vamsi Talla, Mehrdad Hesar, Bryce Kellogg, Ali Najafi, Joshua R Smith, and Shyamnath Gollakota. 2017. Lora backscatter: Enabling the vision of ubiquitous connectivity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 105.

- [56] Viktor Toldov, Laurent Clavier, and Nathalie Mitton. 2018. Multi-channel Distributed MAC protocol for WSN-based wildlife monitoring. In *2018 14th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*. IEEE, 1–8.
- [57] Viktor Toldov, J.P. Meijers, Roman Igual-Perez, Riaan Wolhuter, Nathalie Mitton, and Laurent Clavier. 2016. Performance evaluation of LoRa radio solution for PREDNET wildlife animal tracking project. In *1st Annual LPWAN Conference*.
- [58] Ambuj Varshney, Oliver Harms, Carlos Pérez-Penichet, Christian Rohner, Frederik Hermans, and Thiemo Voigt. 2017. LoRea: A Backscatter architecture that achieves a long communication range. In *The 15th ACM Conference on Embedded Network Sensor Systems (SenSys)*. ACM, 18.
- [59] Veniam. 2016. Creating the World’s Largest Network of Connected Vehicles for Smart Cities. Retrieved Jun 3, 2019 from https://www.worldwifidday.com/wp-content/uploads/2016/05/3.-PortoCaseStudy_Letter_2016-04-15.pdf
- [60] Bill Zalud. 2013. How Mesh Networks Form the Backbone of Smart Cities. Retrieved Jun 3, 2019 from <https://www.securitymagazine.com/articles/84986-how-mesh-networks-form-the-backbone-of-smart-cities>

Received July 2018; revised March 2009; accepted June 2009