

LMAC: Efficient Carrier-Sense Multiple Access for LoRa

Amalinda Gamage, Jansen Christian Liando, Chaojie Gu, Rui Tan and Mo Li

School of Computer Science and Engineering

Nanyang Technological University

Singapore

{amalinda,cjansen,gucj,tanrui,limo}@ntu.edu.sg

ABSTRACT

Current LoRa networks including those following the LoRaWAN specification use the primitive ALOHA mechanism for media access control due to LoRa's lack of carrier sense capability. From our extensive measurements, the Channel Activity Detection (CAD) feature that is recently introduced to LoRa for energy-efficiently detecting preamble chirps, can also detect payload chirps reliably. This sheds light on an efficient carrier-sense multiple access (CSMA) protocol that we call LMAC for LoRa networks. This paper presents the designs of three advancing versions of LMAC that respectively implements CSMA, balances the communication loads among the channels defined by frequencies and spreading factors based on the end nodes' local information and then additionally the gateway's global information. Experiments on a 50-node lab testbed and a 16-node university deployment show that, compared with ALOHA, LMAC brings up to 2.2× goodput improvement and 2.4× reduction of radio energy per successfully delivered frame. Thus, should the LoRaWAN's ALOHA be replaced with LMAC, network performance boosts can be realized.

CCS CONCEPTS

• Networks → Wireless access points, base stations and infrastructure.

KEYWORDS

LoRa, CSMA, LoRaWAN, LPWAN, Chirp Spread Spectrum (CSS), ALOHA, Wireless, IoT.

ACM Reference Format:

Amalinda Gamage, Jansen Christian Liando, Chaojie Gu, Rui Tan and Mo Li. 2020. LMAC: Efficient Carrier-Sense Multiple Access for LoRa. In *The 26th Annual International Conference on Mobile Computing and Networking (MobiCom '20)*, September 21–25, 2020, London, United Kingdom. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3372224.3419200>

1 INTRODUCTION

Low-power wide-area networks (LPWANs) form an important class of wireless networks for geographically distributed Internet-of-Things (IoT) end devices. LPWANs' long-range communication

capabilities will increase the degree of connectivity of IoT and enable deep penetration of networked intelligence into urban territories (e.g., wide areas, buildings, and underground structures) that have challenged the existing low-power short-range wireless technologies. Among various LPWAN technologies (including NB-IoT and SigFox), LoRaWAN [10], an open data link layer specification based on the LoRa physical layer, offers the advantage of using license-free ISM bands, low-cost end devices, and the independence from managed cellular infrastructures.

Currently, the LoRa networks, including those organized in an *ad hoc* manner or by following the LoRaWAN specification, adopt the ALOHA media access control (MAC) mechanism for each logic communication channel defined by radio frequency and the spreading factor (SF) of the chirp spread spectrum (CSS) modulation. The primitive ALOHA, though enabling simple network implementation, is not competent for scaling with the communication demand. As such, even though each end device conforms to the channel access time requirement (e.g., 0.1% or 1% duty cycle in Europe [15]), the ALOHA-based LoRa networks will have degraded network performance due to massive collisions when the numbers of end devices grow sharply in this era of IoT [5, 6, 8, 17, 22, 23, 25, 30, 34, 35].

The absence of carrier sense capability on off-the-shelf LoRa end devices impeded studies and implementations of the more advanced carrier-sense multiple access (CSMA) schemes. However, a recently introduced feature called Channel Activity Detection (CAD) sheds light on CSMA for LoRa networks. The CAD, which is available on all of the latest LoRa radios (e.g., SX126x and SX127x), is designed for energy-efficient preamble detection. Although it is not for full-fledged carrier sense, our extensive measurement study shows that CAD can also detect the payload chirps with satisfactory performance. Specifically, it achieves more than 95% accuracy in detecting the occupancy of a logic channel due to an ongoing frame transmission.

Most existing CSMA designs from other wireless technologies are based on energy detection on the spectrum. For instance, they adopt received signal strength (RSS) as an indicator of the presence of an on-going transmission[1, 2]. Simply porting them to LoRa networks may yield inefficiency. To achieve efficient CSMA for LoRa networks, the unique features and constraints of LoRa radios need to be considered. For instance, LoRa frame can traverse below the noise floor; Concurrent transmissions with different SFs are encouraged in the same frequency channel. These features render the traditional energy detection based CSMA designs ineffective. Thus, a clean-slate CSMA redesign based on LoRa's physical layer characteristics is needed. To this end, we follow the *progressive systems development* methodology to design an efficient CSMA-based LoRa MAC (LMAC). The advancing versions of LMAC integrate various new designs and features based on the precedent code base to

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.

MobiCom '20, September 21–25, 2020, London, United Kingdom

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7085-1/20/09...\$15.00

<https://doi.org/10.1145/3372224.3419200>

address additional and realistic networking problems. Their details are as follows.

LMAC-1: LMAC-1 implements the basic functionality of CSMA, i.e., to avoid collisions by refraining from transmission when the channel is busy. We implement the Distributed Inter-Frame Space (DIFS) mechanism that performs a fixed number of CADs and a random back-off (BO) mechanism. LMAC-1 shall outperform ALOHA in terms of packet reception ratio (PRR) and network goodput when the communication demand increases.

LMAC-2: LMAC-1 falls short of balancing the communication loads. As a result, the busier logic channels have higher frame loss rates. LMAC-2 aims to balance the loads among the logic channels by allowing LoRa end nodes to select the less crowded logic channels based on their local information. We devise an *indirect channel probing* approach for each end node to update its knowledge regarding the channels' crowdedness based on the CAD results during DIFS and BO processes. Thus, LMAC-2 remains as a fully distributed MAC that aligns well with the LoRaWAN Class A specification. In addition, it does not introduce extra CAD overhead compared with LMAC-1. LMAC-2's channel selection shall equalize the channel utilization and bring forth a performance improvement.

LMAC-3: In LMAC-2, the LoRa end nodes may develop biased local views regarding the channel loads. In LMAC-3, the gateway broadcasts global views of channel loads using periodic beacons. The beaconing renders LMAC-3 consistent with the LoRaWAN Class B specification. Upon receiving the global view, the LoRa end nodes update their local views to guide the channel selection before transmission. With the gateway's assistance, LMAC-3 shall further improve network performance with respect to LMAC-2.

To support system design and evaluation, we construct a testbed consisting of 50 LoRa end nodes and a LoRa gateway in a lab environment. We conduct extensive, comparative evaluation experiments with goodput, PRR, and radio energy consumption per successfully delivered frame as main performance metrics. The results show that, compared with ALOHA which yields sub-50% PRR, all LMAC versions maintain sup-90% PRR when the communication demand increases. With respect to ALOHA, LMAC-1, LMAC-2, and LMAC-3 respectively bring 1.5 \times , 1.9 \times , and 2.2 \times goodput improvements, and 2.08 \times , 2.37 \times , 2.38 \times reductions in radio energy per successfully delivered frame. In addition, LMAC-3 achieves the most balanced channel loads. The experiments on a deployment of 16 LoRa end nodes in a university area also show significant performance improvements brought by LMAC in comparison to ALOHA.

The contributions of this paper are summarized as follows. **First**, we study the recently introduced LoRa CAD and show its full-fledged channel-selective carrier sense capability. **Second**, we design advancing versions of LMAC to address LoRa's unique features and various realistic networking problems. The fully distributed LMAC-2 and the gateway-assisted LMAC-3 are the final yields of this work that align well with the LoRaWAN Class A and Class B specifications, respectively. **Third**, our testbed experiments show that, compared with ALOHA, LMAC brings doubled goodput and halved radio energy consumption under saturated communication demands. Additionally, we make available an LMAC integrated LoRaWAN codebase in [21].




Figure 1: CAD results in SF12 and SF11 when there is an on-going frame in SF12. Spectrogram shows the frame chirps.

The rest of this paper is organized as follows. §2 presents a LoRa primer. §3 investigates CAD. §4 describes the testbed. §5 designs LMAC. §6 and §7 presents testbed experiments and simulation results, respectively. §8 reviews related work. §9 concludes this paper.

2 LORA PRIMER

In this section, we present a background on LoRa physical layer, LoRaWAN specification, and CAD.

LoRa and LoRaWAN. A typical LoRa network consists of a number of geographically distributed end devices that transmit/receive data to/from one or more gateways. The radio frequency (RF) spectrum allotted to LoRa is also shared with a number of other physical layer protocols. LoRa physical layer divides the spectrum into multiple frequency channels (CHs), where each CH supports both *uplink* and *downlink* communications. LoRa employs CSS modulation, in which each *chirp* linearly sweeps a mandatory predefined bandwidth BW from the minimum frequency f_{min} to the maximum frequency f_{max} of the CH. CSS represents data by the initial frequency of a chirp.

A LoRa frame has three parts: a preamble of eight up chirps, a Start Frame Delimiter (SFD) of $2\frac{1}{4}$ down chirps, and a payload of multiple data chirps. Each data chirp lasts exactly a duration of T_{sym} milliseconds (ms). A data chirp starting from an arbitrary frequency f_x represents an RF symbol out of 2^{SF-1} total symbols where SF is the spreading factor. The chirp starting from f_x linearly increases in frequency over time, reaches f_{max} , wraps around, and resumes its journey from f_{min} back to f_x . As such, each data chirp sweeps the entire bandwidth BW . Demodulation is performed at the gateway by relating the FFT bins of the multiplication of the payload with a series of down chirps.

LoRa physical layer allows concurrent transmissions in the same CH by using the distinct slopes of the concurrent chirps, i.e., SFs. This feature enhances a LoRa end device's ability to co-exist among many end devices that use the same CH. LoRa supports six SFs, from SF7 to SF12. In the rest of this paper, we use the notation CH/SF to denote a *logic channel* in which the simultaneous transmissions collide.

LoRaWAN is a data link layer specification based on LoRa. It defines three classes, namely, Class A (II end devices), Class B (eacon), and Class C (ontinuous listening). While the three classes have distinct characteristics in response times and energy expenditures, they are all based on the primitive ALOHA MAC. Although the




Figure 2: CAD on various types of CSS signals. Top row shows the CSS signals; middle row shows the generated up chirps; bottom row shows the cross-correlation results. (a) CAD on a preamble chirp; (b) CAD on parts of two consecutive preamble chirps; (c) CAD on a data chirp; (d) CAD on parts of two consecutive data chirps; (e) CAD on chirps in a different SF.

requirement on the channel access time (e.g., 0.1% and 1% in Europe) mitigates channel contention, the ALOHA-based LoRaWAN cannot sustain the growth of end devices. Besides MAC, LoRaWAN also includes other link layer functions such as adaptive data rate and payload encryption.

Channel Activity Detection (CAD). LoRa recently included CAD as a power-optimized mechanism to detect preamble chirps preceding full-fledged payload demodulation. This allows a LoRa node to energy-efficiently detect an incoming frame without resolving to power-hungry continuous RX mode. Some details of CAD are as follows. Prior to responding to a CAD request, the LoRa radio should be set on the desired CH/SF. A single CAD lasts $T_{sym} + \frac{32}{BW}$ ms, during which the radio takes the following steps. First, a lock in the phase-locked loop is achieved to tune the radio to the desired CH. The radio then switches to the receiving mode for T_{sym} ms. Second, the radio switches to processing mode to search for strong cross-correlation between the received I/Q samples and a locally generated up chirp of the same SF. The cross-correlation computation takes $\frac{32}{BW}$ ms. Lastly, a CadDone interrupt is generated. If the cross-correlation is high, a CadDetected interrupt is also generated.

3 ENABLING LORA CARRIER SENSE

Carrier sense capability is desirable for advancing LoRa. Two factors, however, challenge effective LoRa carrier sense. First, LoRaWAN encourages concurrent transmissions in a single CH using different SFs. Second, most CSS-based modulation methods including LoRa allow signals to traverse below the noise floor. For instance, a LoRa receiver’s sensitivity is 10^5 times higher than that of Wi-Fi [11, 31]. The above two factors render the traditional RSS-based channel occupancy detection futile. This section extends the understanding on LoRa CAD beyond the technical documentation [4, 28]. It also evaluates the effectiveness of using CAD for carrier sense.

To use CAD for carrier sense, we need to consider three factors: 1) How reliable is CAD? 2) Could CAD be performed on targeted SFs while the CH is also used by transmissions in others SFs? 3) How frequently can consecutive CADs be performed? To answer

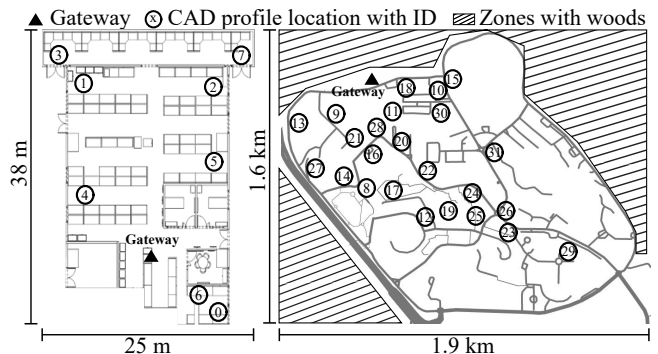


Figure 3: Indoor 0-7 & outdoor 8-31 locations where CADs are performed.

these questions, we conduct an experiment with three LoRa nodes and a software-defined radio (SDR) setup. In the experiment, a node transmits an SF12 frame on CH1 (868.1 MHz, 125 kHz BW). Two other nodes tune to CH1 and continuously perform CAD in SF11 and SF12, respectively. Fig. 1 shows a portion of the CH1 spectrogram recorded by the SDR and the two nodes’ CAD results. We have the following two observations. First, CAD can detect data symbols. Second, CAD strictly adheres to the specified SF and does not generate false positives for other transmissions of different SFs in the same channel.

To develop a better understanding on the behavior of CAD, we conduct more experiments on the cross-correlation process of CAD. Specifically, we compute the cross-correlation using two SF9 and SF12 frames recorded by the SDR. First, we test the case in which the incoming SF12 preamble chirp aligns ideally with the locally generated SF12 up chirp. The cross-correlation result shown in Fig. 2a indicates a clear spike signifying a detection. We then introduce a time offset of $\frac{T_{sym}}{2}$ ms, as shown in Fig. 2b on the incoming SF12 preamble chirps. As a result, the locally generated up chirp covers portions of two consecutive incoming chirps. The cross-correlation produces a clear spike as well. Next, we compute cross-correlation

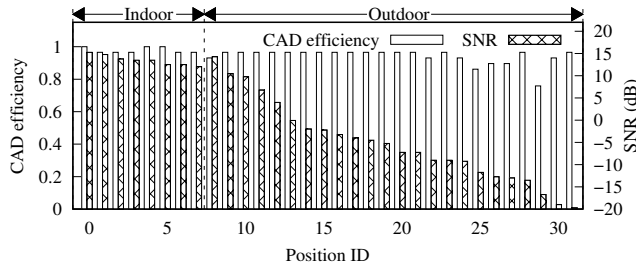


Figure 4: CAD efficiency and SNR at various locations.




Figure 5: Current draw of SX1276 when it is idle, performing CAD, and transmitting in SF7.

for an ideally aligned incoming data chirp which again results a spike as shown in Fig. 2c. Afterwards, we compute cross-correlation for a data chirp with a time offset of $\frac{T_{sym}}{2}$ ms, and observe a spike too as shown in Fig. 2d. Note that the above simulated offsets can be naturally incurred when the CAD process starts at an arbitrary point of the incoming frame. This can be observed from Fig. 1, i.e., a successful CAD does not necessarily need alignment with the incoming symbol. Lastly, we investigate how cross-correlation behaves for chirps in different SF. In Fig. 2e, CAD is performed between incoming SF9 chirps and locally generated SF12 chirps. It yields no peaks.

Efficiency of CAD indoors and outdoors. LPWANs are typically deployed over large geographical areas. We also investigate the efficacy of CAD in an outdoor area. We deploy two LoRa gateways, one indoor and the other outdoor, and use a mobile LoRa node to perform CAD at multiple indoor and outdoor locations. At each location, we record the signal-to-noise ratio (SNR) reported by the LoRa node when successfully receiving a frame and the CAD efficiency that is the ratio of successful CADs among all CADs performed. Fig. 3 illustrates the experiment scenarios of (a) eight indoor locations within an office of 800 m² and (b) 23 outdoor locations scattered across a university area of 3.5 km². Fig. 4 shows the SNR and CAD efficiency at all visited locations. The median CAD efficiency values are 98% and 94% indoors and outdoors respectively with SNR ranging from -15 dB to +15 dB.

CAD energy footprint. Since CSMA introduces an additional energy overhead for each frame transmission, we profile the power consumption of CAD. To capture the power trace, we connect a power monitor [29] in series to the power supply pins of the SX1276 radio chip. Fig. 5 depicts the incurred current of a single CAD in comparison with those during the idle and transmission states. CAD

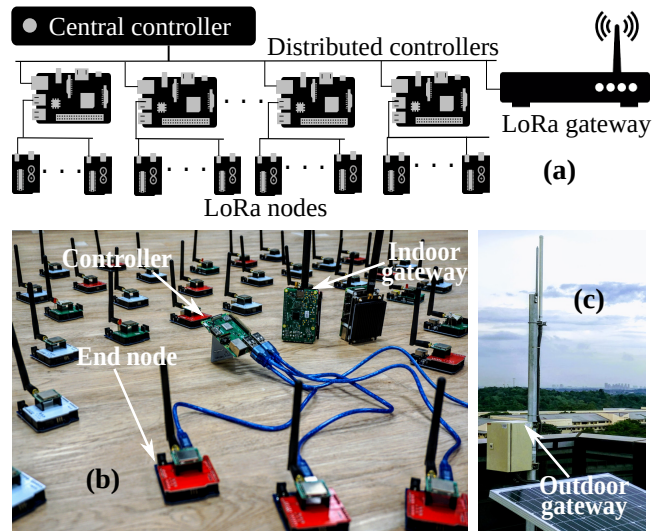


Figure 6: (a) Testbed architecture. (b) Indoor testbed devices. (c) Outdoor gateway.

draws 0.03 W, which is 4.4× of the idle power and 11× less than the transmission power. A CAD in SF7 lasts for 1.2 ms only, consuming 0.029 mJ only. Under SF12, it lasts for 31.6 ms and consumes 1.1 mJ.

Summary. The above experiments give the following key observations. (1) CAD can be used as an effective and reliable carrier sense method; (2) CAD is SF-selective, i.e., it can detect transmissions in a specified SF while remaining insensitive to other transmissions in different SFs; (3) CAD is an energy-efficient operation compared with the energy consumption of data transmissions; (4) CAD works indoors and outdoors with high detection efficiency.

4 LORA TESTBED DESIGN

To facilitate the design and evaluation of LMAC, we build a scalable LoRa testbed. This section presents the testbed design and experiment workflow which has been used for both indoor and outdoor experiments.

Testbed design. We aim to meet the following design requirements. First, the testbed should support a sparse deployment, i.e., across a university. Second, it should support convenient experiment configuration via a central controller. Third, it should provide a bidirectional backhaul link to disseminate configurations and receive experiment logs from the distributed LoRa nodes. Meeting these requirements imposes implementation challenges. For instance, it is non-trivial to synchronously control many LoRa nodes.

Connecting each LoRa node to the *central controller* is not feasible as the LoRa nodes do not have wideband radios to support backhaul links. As such, we employ extensible units, each a cluster of LoRa nodes, and call them *distributed controllers*. The central controller and the distributed controllers are Raspberry Pi single-board computers that perform two coordinated roles. The distributed controllers execute and if necessary, relay control commands to LoRa nodes, thereby creating wideband out-of-band control links between the central controller and LoRa nodes.

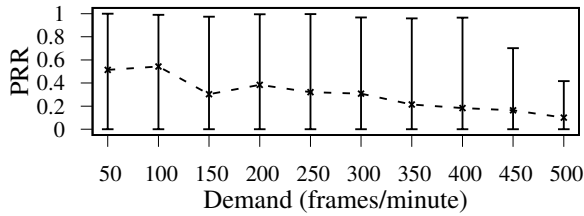


Figure 7: PRR vs. demand under ALOHA. Error bar shows the max, min, and mean PRR among 50 nodes.

Each testbed LoRa node is assembled in house by integrating an SX1276 LoRa transceiver with an Arduino Atmega328p MCU via a bridge board having a capacitor to support transient power requirements unsupported by USB power delivery [24]. Communications between LoRa nodes and a distributed controller (e.g., protocol firmware downloading and experiment log uploading) happens via an emulated serial interface. Fig. 6a illustrates the connectivity between each component of the testbed; Fig. 6b shows a shot of testbed hardware; Fig. 6c shows the outdoor gateway. To address the requirement of supporting geographically distributed nodes, testbed devices communicate through a network tunnel [12] from their local Ethernet.

Experiment workflow. An experiment includes the following steps: establishing connectivity among all components, loading a MAC protocol into LoRa nodes, parsing experiment configurations, executing the experiments, and finally collating logs. In what follows, we provide some important details. At the beginning of an experiment, the central controller distributes a pre-compiled binary MCU program implementing a MAC protocol and two parameters (demand and experiment duration) to the distributed controller. The demand is the requested number of LoRa frames to be transmitted within the specified experiment duration. The extent that the requested demand is met reflects the network performance of the tested protocol. Varying the demand creates different levels of contention. After a mutual confirmation process to ensure that devices are ready, the central controller issues a start command to launch an experiment. Once the experiment duration expires, the distributed controllers report the following experiment logs back to the central controller: 1) LoRa node ID, 2) demand, 3) transmitted packets, 4) CADs performed on each CH/SF, and 5) the count of repeated experiments. In our experiments, all LoRa frames contain 16-byte payloads with the following fields to facilitate performance counting: 1) node ID, 2) packet counter, 3) demand, 4) the number of padding bits.

5 DESIGN OF LMAC

This section presents the progressive design of LMAC to exploit CAD for CSMA-based MAC in LoRa networks. Under ALOHA, the end nodes access the medium abruptly regardless of the channel status, which results in excessive collisions when the communication demand is high. Fig. 7 shows an experiment result in which the average packet reception ratio (PRR) of ALOHA drops from 0.5 to 0.2, when the demand increases. The first version of LMAC, LMAC-1, uses CAD-based carrier sense to avoid the collisions.




Figure 8: Following LMAC-1, node-B detects a collision when node-A is transmitting in the same CH/SF.

5.1 LMAC-1

LMAC-1 is largely based on the Distributed Coordination Function (DCF) that serves as one of the channel access methods of the IEEE 802.11. However, as LoRa differs significantly from 802.11 physical layer, changes are required to adapt DCF to LoRa. A key difference between 802.11 and LoRa is that LoRa does not have a feedback channel and thus cannot support the acknowledgment mechanism of DCF.

In LMAC-1, the CH/SF is chosen randomly for each transmission. An LMAC-1 node transmits a frame after the selected CH/SF is confirmed to be idle. To this end, LMAC-1 adopts the Distributed Inter-Frame Space (DIFS) mechanism from DCF with a fixed number of CADs per DIFS slot. When a DIFS slot completes with all CADs reporting idle channel, the MCU generates a random back-off (BO) value (N_{BO}) and enters the BO phase immediately. Then, the node decrements N_{BO} per each CAD reporting idle channel. The randomized initial value for N_{BO} reduces the possibility where two or more frames collide should the DIFS processes happen to start at the same time. Upon N_{BO} reaching zero, the node transmits the frame. Since a node continues checking the availability of the channel during the BO phase, an occupied channel will reset the node to the DIFS state. Should that occurs, the node needs to wait for an idle channel again and resumes counting down N_{BO} . Note that the initial N_{BO} value is not regenerated until the frame is transmitted. In the rest of this paper, we refer to CADs that report a busy channel as *failed* CADs and vice versa. A failed CAD operation within DIFS or BO will render a failed DIFS or BO slot respectively. Fig. 8 illustrates the process in which two nodes contend for the same CH/SF under LMAC-1. Node B detects an ongoing transmission during DIFS and continues performing DIFS until the channel is idle. Upon detecting an idle channel, node B enters the BO state with $N_{BO} = 3$ and starts decrementing N_{BO} with each successful CAD. Once N_{BO} reaches zero, node B transmits the frame. In the implementation of LMAC-1, we set the range for the initial value of N_{BO} to be $[4, 64]$ and 12 CAD operations per DIFS slot.

LMAC-1 focuses on avoiding collisions. From the evaluation results that will be presented in §6, LMAC-1 improves the PRR significantly over ALOHA. In what follows, we present an experiment result that shows the imperfection of LMAC-1 and motivates LMAC-2. Fig. 9 presents the numbers of frames transmitted by all the end devices (labeled TX) and received by the gateway (labeled RX) in each CH/SF during an experiment. The difference between the TX and RX bars in a certain CH represents the number of lost frames. From the figure, an important observation is that, the busy channels (i.e., those with high bars) have more lost frames. The




Figure 9: Numbers of frames transmitted by all the nodes (TX) and received by the gateway (RX) under LMAC-1. More utilized channels with higher bars suffer from more frame losses, i.e., larger differences between TX and RX bars.

cause of the frame losses is explained as follows. The gateway radio chip SX1301 incorporates eight decoders that can be configured to demodulate incoming frames concurrently. Each decoder must be configured to a single CH but listens to all SFs. However, all decoders share an I/Q buffer. Should the decoders lag behind in decoding incoming I/Q data, a buffer overflow will occur, resulting in frame loss. This is because a decoder is tied to a specific frequency CH and a large number of frames arriving in that CH result in that a decoder lags behind compared to others. We refer interested reader to [20] for a thorough analysis of concurrent frame reception. The design of the SX1301 chip and the results shown in Fig. 9 imply that if the loads of the logic channels are equalized, the communication performance can be further improved.

To this end, the nodes can choose underutilized or idle channels instead of contending for a highly utilized channel. LMAC-2 aims to implement this feature.

5.2 LMAC-2

To select an underutilized channel, a node needs to track the utilization levels of all channels. Since continuous channel probing is time consuming, LMAC-2 leverages information from failed DIFS and BO operations to enable *indirect channel probing*. Thus, each LMAC-2 node maintains historical information regarding the crowdedness of past sensed CH/SF combinations. Later, the LMAC-2 node consults this information to shift to a better CH/SF when a channel is sensed busy during DIFS or BO. Doing so leads to more equalized channel loads and thus, better network performance. In addition, from the perspective of a node, the average number of CADs per transmission and the related time/energy expenditure can be reduced since busy channels are avoided. The key differences between LMAC-1 and LMAC-2 are as follows. In LMAC-1, a node fixates to a random CH/SF until a frame is transmitted. In contrast, an LMAC-2 node uses historical information to dynamically select the best CH/SF. The details of LMAC-2 are presented as follows.

Each LMAC-2 node maintains a *channel occupancy matrix* (Γ) of size $N_{CH} \times N_{SF}$, where N_{CH} and N_{SF} respectively represent the numbers of CHs and SFs. Each element in Γ , denoted by $\gamma_{CH,SF}$ represents the historical utilization rate for a given CH/SF. The $\gamma_{CH,SF}$ is updated for the $t + 1^{th}$ time whenever a node has a failed CAD during DIFS or BO. We adopt exponential averaging for the updating: $\gamma(t + 1) = \alpha \left(\frac{CAD_{busy}}{CAD_{total}} \right) + (1 - \alpha)\gamma(t)$, where α is the learning rate, CAD_{busy} and CAD_{total} are the number of CADs performed when the channel is busy and the total number of CADs

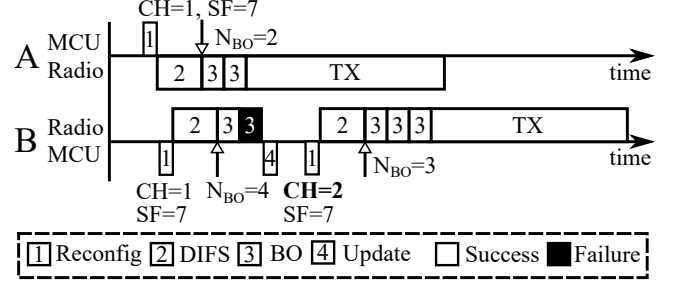


Figure 10: In LMAC-2, upon detecting that the selected CH/SF is unavailable, the node will search for the next CH/SF and perform DIFS on the new CH/SF.




Figure 11: Channel occupancy matrices (Γ) maintained by node 24, node 17, node 12, and gateway in a single experiment. A hatched pattern represents a null value due to no CADs performed in the corresponding CH/SF from the start of the experiment.

performed respectively. The above exponential averaging is used for adapting to channel dynamics whilst maintaining historical information. As LoRa channels are highly dynamic, in our experiments, we set $\alpha = 0.8$ to ensure that higher importance is placed on the latest sensed channels.

We now describe how an LMAC-2 node uses Γ to select the next best CH/SF. A real-world LoRa deployment is subject to dynamic channel conditions. A channel that is seemingly busy may later become free. Therefore, simply choosing the least busy CH/SF according to the γ value may not always reap optimal network performance. As such, we randomize the choice. Specifically, we choose the best, the second best, and the third best CH/SF in terms of γ value with probabilities of 0.5, 0.3, and 0.2, respectively. This allows a forgiving approach for channels that were once crowded.

The next SF selected is set higher or equal to that configured during deployment to ensure the frame robustness.

We use an example shown in Fig. 10 to illustrate the interaction between two LMAC-2 nodes that have selected the same CH/SF. Since node A completes BO earlier than node B, it proceeds with the transmission and node B reports a failed BO during node A's transmission. Thus, node B updates the occupancy matrix element corresponding to the current CH/SF, and re-configures the radio to the next CH/SF. Since the newly selected CH/SF is different from that used by node A, node B can complete DIFS and BO successfully and transmit the frame. If following LMAC-1, node B will wait till node A completes transmission.

As indirect channel probing is opportunistic, the channel occupancy matrices of LMAC-2 nodes are different. We illustrate this in Fig. 11. Specifically, Figs. 11a, 11b, and 11c visualize the Γ

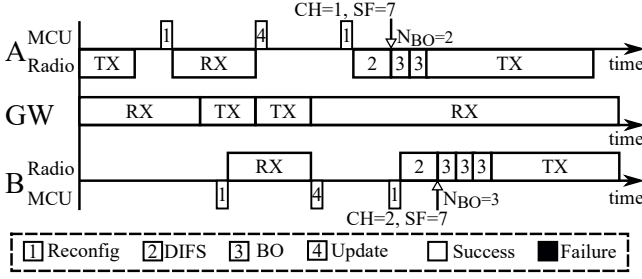


Figure 12: Overview of LMAC-3 protocol.

matrices maintained by node 24, node 17, and node 12 during an experiment. They are different. This difference leads to different channel selections of LMAC-2 nodes. LMAC-2 has the following two issues that motivate LMAC-3. First, some elements in the Γ matrices remain null because no CADs have been performed on those corresponding CH/SF yet. Channels with null γ values may be busy. Without further information, avoiding these channels is wise but also losses opportunities in cases they are idle. Second, certain elements in Γ can be outdated due to lack of proactive channel probing. Addressing above two issues can advance LMAC to further reduce contention. To this end, we can exploit the global view of the gateway to complement the local views of LMAC-2 nodes. Fig. 11d shows the global channel occupancy matrix maintained by the gateway, which reflects the loads of all CH/SF combinations in real time. It does not suffer from null values and outdated issues. LMAC-3 is developed to use this global channel occupancy matrix.

5.3 LMAC-3

In LMAC-3, a gateway transmits a periodic beacon that contains the global channel occupancy matrix (Ψ) to the nodes in the network. Promptly, each receiving node follows a merging process to supplement its local occupancy matrix Γ with the gateway's global occupancy matrix Ψ . The Ψ is an $N_{CH} \times N_{SF}$ matrix where each element $\Psi_{CH,SF}$ represents the contention level in a given CH/SF.

Several challenges however need to be addressed for large-scale LoRa deployments to benefit from the global views of gateways. First, the process of disseminating a gateway's global view across the network should not be detrimental on the gateway's reception capability. Second, a LoRa node may receive overlapping coverage from two or more gateways. Due to network dynamics, a given CH/SF will face dissimilar contention across gateways.

Most commercial LoRa gateways only equip a single antenna and do not support duplex communications. As such, responding to each uplink frame with the Ψ matrix will reduce reception performance. Thus, to disseminate a gateway's global view, LMAC-3 employs time-distributed beacons via a predefined feedback channel similar to LoRaWAN Class B. Fig. 12 shows an example of this process where a gateway disseminates its Ψ to two nodes. Under LMAC-3, a LoRa gateway mostly stays in receiving mode and only switches to transmitting mode when it broadcasts beacons containing Ψ . Note that the nodes may miss a single beacon due to clock drifts. Thus, the beacon is transmitted twice on the feedback channel to increase the probability of reception. LMAC-3 does not require tight clock synchronization across nodes but coarse synchronization. Upon

receiving Ψ , the nodes will perform the matrix merging that will be explained shortly and continue to transmit pending frames while awaiting the next beacon period. The Ψ matrix and the gateway ID is carried by a frame payload of 49 bytes where, an element $\Psi_{CH,SF}$ is represented by a 8-bit non-negative integer with 254 representing the highest load.

The $\Psi_{CH,SF}$ is formally defined as the ratio of the chirps that utilized the channel ($S_{CH,SF}^{RX}$) over the maximum possible during a beaconing period (T_{beacon}). The maximum number of chirps during a T_{beacon} is the multiplication of the chirp rate ($S_{CH,SF}$) and T_{beacon} . The chirp rate is defined as the reciprocal of the symbol duration in the considered SF. Following above, $\Psi_{CH,SF} = \frac{S_{CH,SF}^{RX}}{S_{CH,SF} \times T_{beacon}}$. To compute $S_{CH,SF}^{RX}$ for a given CH/SF, we count the number of chirps in the totally rx received frames in the CH/SF as follows: $S_{CH,SF}^{RX} = \sum_{i=0}^{rx} (n^{sfd} + n_i^{pre} + n_i^{pl})$, where n^{sfd} is the number of the SFD chirps in a frame, n_i^{pre} and n_i^{pl} are the numbers of preamble chirps and payload data chirps in the i^{th} received frame [3].

We now discuss how to merge Γ and Ψ . The simplistic approach of using Ψ solely to guide the channel selection is problematic, because the nodes will likely choose the same channel causing channel occupancy imbalance and contention. Our merging process adopts an element-wise weighted sum of Γ and Ψ , with a large weight for Γ and a small weight for Ψ . Thus, the merging result preserves the diversity of the nodes' observations while capturing the global observations by the gateway.

To address the multiple gateway scenario, a gateway beacon integrates a single byte ID which aids in identifying different gateways. The $\Psi_{CH,SF}$ which has normalized channel contentions across gateways allows for simple merging of multiple Ψ matrices at node level. The above enables LMAC-3 to be scalable.

Based on our empirical results under different environment conditions for varying contentions, we choose 0.8 and 0.4 respectively as weights for Γ and Ψ . The experiment results presented in §6 suggest that the determined weights provide high performance across various conditions and therefore can be used as is by new networks.

6 EVALUATION

We conduct indoor and outdoor experiments using the testbed described in §4 to evaluate LMAC. We use PRR, goodput, and average energy consumption per successful reception as the performance metrics. We also consider protocol fairness in terms of the number of CADs performed by each node and distribution of channel loads.

6.1 Indoor Experiments

We conduct indoor experiments with co-located nodes in a lab environment, which allows us to better control the process and cover an extensive range of parameter settings. The experiments use 50 LoRa nodes. There are total of 16 CH/SF combinations used in the experiments (eight CHs and two SFs). In the experiments, we adopt the following default settings: 16-byte payload, 10-chirp preamble, and 2dB transmission power. Each experiment lasts one minute and is conducted in different times of the days over two months.

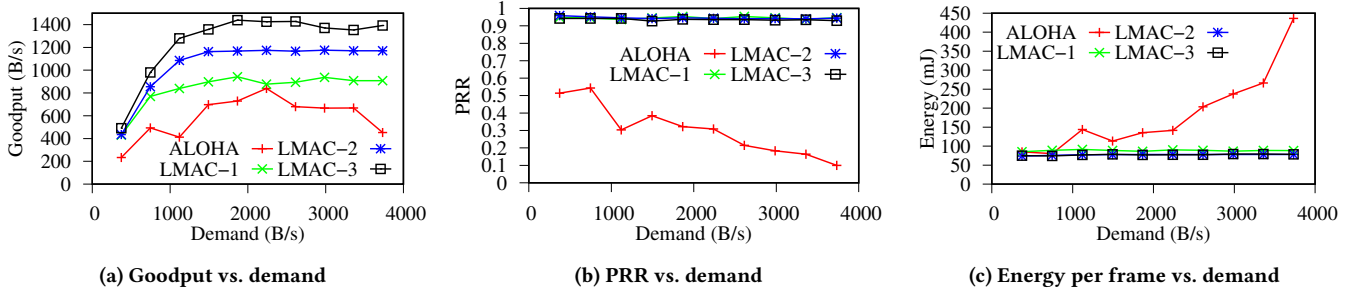


Figure 13: Network performance comparison among ALOHA, LMAC-1, LMAC-2, and LMAC-3. LMACs maintain about 90% network PRRs, achieve higher goodputs, and reduce energy consumption per frame reception.




Figure 14: CDFs of the per-node network performance under ALOHA, and LMAC-1/2/3 when the network demand is 2,600 B/s. Most LMAC nodes outperform ALOHA nodes except for a small set of ALOHA nodes. The steeper CDFs of LMACs suggest that the performance metrics are fairer among the nodes.

Overall network performance. We evaluate the goodput, PRR, and node energy consumption when network demand varies under the four MAC protocols. Network goodput is the amount of meaningful data delivered per second. From Fig. 13a, we can see that compared with ALOHA, LMAC-1/2/3 improve the goodput by 1.52 \times , 1.87 \times , and 2.21 \times , respectively. Thus, LMACs can deliver data efficiently owing to the collision avoidance. Moreover, the performance increments of LMACs suggest the effectiveness of utilizing channel diversity and CH/SF selection. ALOHA goodput rises as the network demand increases but reduces once network demand is beyond 2500 B/s. This is consistent with the existing understanding of ALOHA performance. The goodputs of LMACs saturate and remain flat once the demand is beyond 2000 B/s. The saturation of LMACs goodput suggests the contention in channel utilization.

From Fig. 13b, the PRRs of LMACs are maintained above 90%, while that of ALOHA decreases with demand. The decline is due to increasing collisions. Differently, LMACs maintain high PRRs through avoiding collisions and busy CH/SF.

Fig. 13c shows the average energy consumed by a node per successfully delivered frame under the four MAC protocols. Compared with ALOHA, LMAC-1/2/3 reduce the energy by 2.08 \times , 2.37 \times , and 2.38 \times , respectively. Fig. 13c also shows that LMAC-2 and LMAC-3 consume similarly low energy, whereas LMAC-1 consumes slightly higher energy than LMAC-2/3. It is because LMAC-1 nodes continuously perform DIFS to contend for the channel. Differently, LMAC-2/3 perform less CADs due to wise CH/SF selection. In the

experiment logs we find that LMAC-3 performs 1.92% less CADs than LMAC-2, whereas LMAC-2 performs 7.81% less CADs than LMAC-1. The goodput gain of LMAC-3 over LMAC-2 is due to the time saved from less CADs performed. Specifically, LMAC-3 nodes can utilize the saved time for frame transmissions, while the PRRs of LMAC-2/3 are similar.

Overall, LMAC outperforms ALOHA through collision avoidance and achieves up to 2.21 \times goodput improvement and 2.38 \times reduction in nodes energy consumption, whilst maintaining above 90% PRR.

Per-node performance. A closer look at the results from a single experiment will provide more insights into the advantages of LMAC over ALOHA. Fig. 14 shows the cumulative distribution functions (CDFs) of the per-node network performance metrics. From Fig. 14a, majority of ALOHA nodes suffer low goodputs that are smaller than 5 B/s, while the remaining ALOHA nodes achieve goodputs scattered from 28 B/s to 74 B/s. The standard deviation of ALOHA’s per-node goodput is 26.79 B/s. In contrast, under LMACs, the CDFs do not exhibit undesirable long tails; the per-node goodputs are mostly within the range of 20 B/s to 60 B/s, with a standard deviation of 14.89 B/s. The above results suggest that, compared with ALOHA, LMACs achieve better balance among the nodes in utilizing the communication medium.

Fig. 14b shows the CDFs of the per-node PRRs. LMAC-2/3 achieve at least 80% PRR. The mean PRR values of all LMACs are above 90%. In contrast, about 40% ALOHA nodes suffer zero PRR, because the frames are lost due to collisions. The remaining nodes have scattered PRRs from 0 to 95%. Experiment logs suggest that ALOHA




Figure 15: Under LMAC-2/3, both PRR and frequency channel load have limited variations. LMAC-1 may have lower PRRs. Under ALOHA, the frequency channel loads are scattered and the PRRs are low.

nodes with high PRRs dominate the CH/SF and hinder the reception of other nodes.

Fig. 14c shows the CDFs of the per-node energy consumption. Up to 30% ALOHA nodes cannot deliver a single packet. Thus, they are not accounted in the CDF. Another 40% ALOHA nodes consume excessive per-frame energy due to collisions. LMAC nodes are rarely affected by frame losses. The total energy consumed for channel probing trades for the exemption of futile energy consumption for transmitting lost frames. Overall, all LMAC nodes outperform 85% ALOHA nodes in terms of energy consumption.

In summary, a closer look into individual node performance reveals significant reduction in energy consumption per node which is hidden in the earlier results due to averaging. Moreover, under LMACs, the goodput and PRRs of each node are more evenly distributed among nodes, while ALOHA nodes are clearly divided into the two extremes with a small number of exceptional nodes.

Balance and fairness. We evaluate the loads of all frequency channels under various MAC protocols. Since each channel supports multiple SFs, the frequency channel load (Ψ_{CH}) needs to be aggregated properly from the logic channel loads ($\Psi_{CH,SF}$) that are estimated by the gateway in the experiments. By averaging from SF7 to SF12 in each channel, the channel load for each frequency channel is obtained.

Fig. 15 shows the scatter plot of PRR versus channel load. The area of $PRR \in [0.8, 1]$, $\Psi_{CH} \in [0, 0.1]$ is zoomed in to present a detailed view of LMAC performances. PRR of ALOHA drops quickly with the frequency channel load. The result of ALOHA shown in Fig. 15 reinforces our understanding on the disadvantages of ALOHA’s completely uncoordinated transmissions. In contrast, the frequency channel loads of LMAC-1 are maintained within 0.1. This low load helps maintain satisfactory PRRs. It shows that LMAC-1 effectively avoids frequency channel saturation. More importantly, the channel selection mechanisms make LMAC-2 and LMAC-3 operating within the region with frequency channel load smaller than 0.1 and PRR within [80%, 100%]

Fig. 16 shows the sorted average numbers of CADs per frame performed by LMAC-2 and LMAC-3. LMAC-2 performs more CADs and yields a steeper slope which means higher variation. The additional channel information provided by the gateway helps LMAC-3 nodes to avoid known congested channels and perform less CADs.

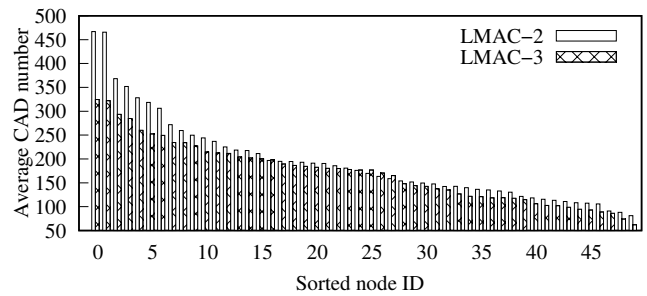


Figure 16: Average CADs performed per frame by each node under LMAC-2 and LMAC-3 when the demand is 500 frames a minute.

Many CADs are performed for every frame due to the DIFS design. Recall that each DIFS slot contains 12 CADs. Moreover, CADs are also performed during the decrement process of N_{BO} .

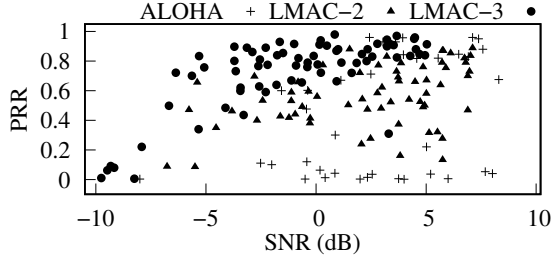
With the above experiments, LMACs yield a controlled and a fairer frequency channel utilization. This allows node performances to be more evenly distributed and ensure data delivery for each node. Moreover, the overhead of LMAC-2/3 is shown to be spread evenly among nodes where most nodes perform less than 250 CADs or 21 DIFS per transmission within a high contention network. The fewer CADs with high reception probability save radio energy for each successful frame delivery.

The latency introduced by LMAC is largely bound to network’s contention levels. LMACs, under high contention levels, will incur higher latency due to the wait for free channels. ALOHA on the other hand is not subject to this latency but suffers from much higher frame loss rate due to collisions, which translates to much higher delay if retransmissions are used to ensure the same level of PRR.


6.2 Outdoor Experiments

To investigate the feasibility and performance of LMAC in an outdoor environment, a set of experiments were conducted in our university area. The experiment makes use of four distributed controllers and 16 end nodes. To exhibit the advantages of LMAC in a high-density network with low number of end nodes, the outdoor experiments use six CH/SF combinations. This set of experiments focus on LMAC-2 and LMAC-3 as the final yields of this paper. We collect 30 minutes of experiment traces with 10 minutes for all protocols including ALOHA. Different from the indoor experiments, the outdoor experiments were conducted between 10AM to 10PM spanning across seven days.


The results of the experiments are presented in Fig. 17. Each point in the figure represents a node with the averaged SNR when the gateway received the frames. The standard deviation of SNR is between 0 to 3dB. The PRR (Fig. 17a) and goodput (Fig. 17b) of LMAC-2/3 spread across vertically and begin to concentrate downwards as SNR reduces. This shows the nodes with higher SNRs (i.e., well placed or nearby nodes) can deliver their frames more reliably. On the other hand, nodes with lower SNRs (i.e., poorly placed or obstructed nodes) most likely cannot reach the gateway and fail to deliver their frames. In contrast, ALOHA nodes



(a) Average SNR vs. PRR




(b) Average SNR vs. goodput

Figure 17: Performance of nodes in an outdoor environment from the perspective of average SNR of a node.

Figure 18: Cumulative effective traffic as SNR increases.

are scattered in the plots with no clear pattern. Although there are nodes that achieve very high PRRs and goodputs, the majority of nodes have low PRRs and goodputs. The average PRR and goodput of ALOHA are 0.22 and 15.98 B/s respectively, while LMAC-2/3 reports 0.49 and 0.63 average PRRs with 25.17 B/s and 34.83 B/s average goodputs respectively. Notice that ALOHA has less points than LMAC-2/3. This is because more than 40% of ALOHA nodes' frames are lost at the gateway resulting in no SNR values registered for those nodes.

To understand the actual traffic of the network for each protocol, Fig. 18 presents each frame contribution to the effective traffic of the respective protocol. The results show that high SNR frames are more likely to be received by the gateway. Notice the effective traffic of ALOHA only starts to increase when SNR reaches 4dB while LMAC-2 and LMAC-3 effective traffic starts to increase at an earlier SNR. The late increase of ALOHA shows that frames with lower SNR are often drowned by higher SNR frames. Differently, LMAC-2/3 allow low-SNR nodes to transmit only when the channel is assessed to have low collision probability. Moreover, the total


Figure 19: Percentage energy overhead of LMACs versus demand.

effective traffic of ALOHA is lower than those of LMAC-2/3. This coincides with the earlier observation where more than 40% nodes are not accounted for due to no frame reception.

The above results suggest LMAC can be affected by environment dynamics but it provides effective traffic control even for low-SNR nodes. Contrarily, ALOHA nodes are divided into two extremes where one yields good performance while the other suffers from low PRR hence low goodput.

6.3 Energy Overhead

LMAC energy overhead is mainly attributed to CADs performed under *DIFS* and *BO* windows. We derive the percentage energy overhead of LMAC in comparison with a conventional ALOHA transmission.

We review the indoor experiment results based on 50 nodes with 10 different traffic demands in §6.1 and derive the average number of CADs a node needs to perform to send out a frame with the different network contentions, denoted by N_{CADs} . N_{CADs} accounts to average CADs performed during *DIFS* and *BO* windows across multiple SFs. From that, the average energy overhead per frame of LMAC, denoted by E_{LMAC} , can be computed as $E_{LMAC} = E_{CAD} \times N_{CADs}$, where E_{CAD} denotes the average energy consumption per CAD across all SFs. Therefore, the per-frame percentage energy overhead of LMAC is estimated as $\frac{E_{LMAC}}{E_{frame}} \times 100\%$ where E_{frame} denotes the transmission energy for a single LoRa frame averaged across all SFs. We estimate through energy measurements that E_{CAD} and E_{frame} correspond to 0.356 mJ and 0.31 J, respectively. We also find through measurements that N_{CADs} is similar across LMAC-2 and LMAC-3, which is also signified by the per frame energy presented in Fig.13c. Therefore, we compute the energy overhead of LMAC as a single entity for both LMAC-2 and LMAC-3. We then extend the percentage energy overheads of LMAC for varied network demands in Fig. 19. We see an energy overhead of about 20% to 25% which does not significantly grow with the increase of network demands (thus higher contentions).

Simply computing LMAC overhead in comparison with ALOHA frame energy E_{frame} does not fairly represent a real LPWAN network's performance that suffers from frame loss. We factor in the frame loss when considering the energy overhead of goodput, and derive the PRR calibrated energy overhead as expressed in Eq. 1,

$$\left[\frac{E_{LMAC} + E_{rframe(LMAC)}}{E_{rframe(ALOHA)}} - 1 \right] \times 100\% \quad (1)$$

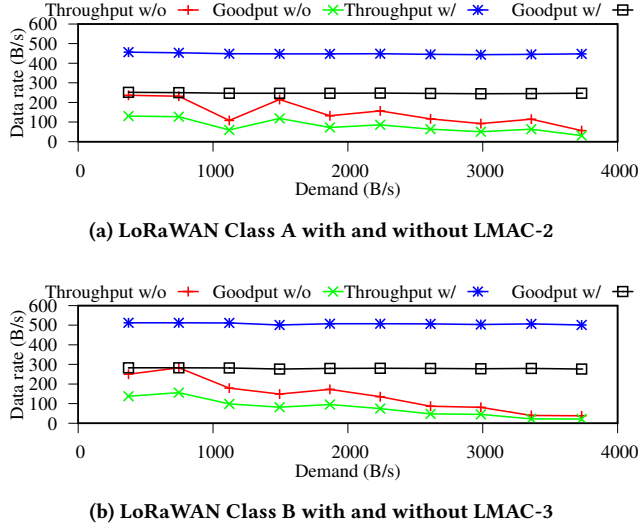


Figure 20: LoRaWAN performance.

where $Er_{frame}(LMAC)$ and $Er_{frame}(ALOHA)$ correspond to the PRR calibrated per frame energy, i.e., $\frac{E_{frame}}{PRR(LMAC)}$ and $\frac{E_{frame}}{PRR(ALOHA)}$, respectively.

The two quantities characterize the energy costs per successfully transmitted frame with LMAC and ALOHA. We present the PRR calibrated energy overhead of LMAC atop ALOHA in Fig. 19 as well. The results show negative overhead and attest that significant energy gains (40% to over 80% energy savings) are achieved with LMAC, especially under high network contentions where ALOHA is subject to much higher frame loss. Fig. 19 suggests that at a light energy cost in CADs, LMAC is able to greatly improve the PRR, which gains energy savings in achieving the same amount of goodput.

7 LORAWAN IMPLICATIONS

The testbed experiments in §6 have shown that LMACs improve the performance of LoRa networks. In this section, we apply LMAC to LoRaWAN which is the prevalent data link layer protocol primarily based on ALOHA. As LoRaWANs are constrained by the channel access time requirements (<1% duty cycle ratio for each end node, etc.), it is difficult to experimentally show the full advantage of LMAC with limited number of devices. Instead, we conduct trace-driven simulations to exhibit the LMAC advantage for a scaled LoRaWAN network. We replace the ALOHA MAC in LoRaWAN Class A with LMAC-2 and LoRaWAN Class B with LMAC-3. LMAC-2 does not require information from the gateway, similar to the autonomy of LoRaWAN Class A; LMAC-3 requires gateway feedback which is provided by the beacon slot of LoRaWAN Class B.

7.1 Data Rates of LMAC-based LoRaWANs

The data traces used are collected from 30 experiments each spanning one minute for each of the three protocols. Thus, there is a 10-minute trace in total for each protocol. Trace frames are added with LoRaWAN headers and mandatory delays (e.g., guard time,

RX slots, etc). With the addition of LoRaWAN overheads, trace durations are recalculated. Once trace duration has been updated, frames received and lost are replayed in the simulation as per trace records. The throughput and goodput of the replayed simulation are then recorded and presented in Fig. 20.

Despite the low data rate, LMAC could still bring 3.06× and 5.93× performance improvements on LoRaWAN Class A with LMAC-2 and LoRaWAN Class B with LMAC-3, respectively. The simulation results suggest the performance of LoRaWAN can be improved by utilizing LMAC as the default MAC. A noticeable difference from the results in §6.1 is the data rate of LoRaWAN decreases due to protocol overheads. As most of the overheads are due to the delays imposed for the two RX slots following a TX, the data rate achieved without the imposed delays will be much higher. However, our simulations still impose such delays to be consistent with the current specification of LoRaWAN.

Since LMAC does not alter any part of the existing LoRaWAN protocol except the MAC mechanism, the integration of LMAC to the LoRaWAN protocol stack could be readily achieved. Moreover, as LoRaWAN imposes channel access restrictions for fairness and LMAC can improve the channel access fairness by avoiding occupied or congested channels, the channel access restrictions can be relaxed or even removed from the LoRaWAN should LMAC be used.

7.2 Network Capacity

Understanding the maximum number of nodes that can be supported by the LMAC-based LoRaWAN is important. This section adopts a numerical analysis method to estimate the number of frames that the LMAC-based LoRaWAN can receive in a given period. Such number can be translated to the maximum number of supported end nodes once the per-node demand is given.

Using the results in previous experiments, the total number of chirps observed by the gateway in a channel (S_{CH}^{RX}) can be extracted within an observation duration (T_{obs}). With the data above, the observable chirps per second in the channel (S_{CH}) can be obtained through $S_{CH}^{RX} \div T_{obs}$. Having S_{CH} , an estimated number of frames a channel can receive (Frm_{CH}^{RX}) can be estimated by simulating different S_{CH}^{RX} values generated from varying frame parameters, i.e., payload size, time period, SF, code rate (CR), etc. With the simulated S_{CH}^{RX} , Frm_{CH}^{RX} is computed by $(S_{CH} \times T_{sim}) \div S_{CH}^{RX}$, where T_{sim} is the simulated time period.

A network analysis using the above method to estimate the network capacity when three parameters (payload size, SF, and CR) vary are presented in Fig. 21. Varying payload size exhibit an exponential decay in the number of frames received shows that the payload size plays a vital role in network capacity. The network capacity declines exponentially as the payload size increases due to frame overhead. Frames with small payload size carry a larger overhead to payload ratio than frames with large payload size. Varying SF also exhibits similar decline as varying payload. Such decline is due to the non-linear increase in number of chirps per frame as SF increases. In contrast, varying CR decline is linear. This linearity reflects the linear increase in each CR setting from CR4/5 to CR4/8 where each increase in setting is an additional bit of redundancy added to every four bits of the frame. The results

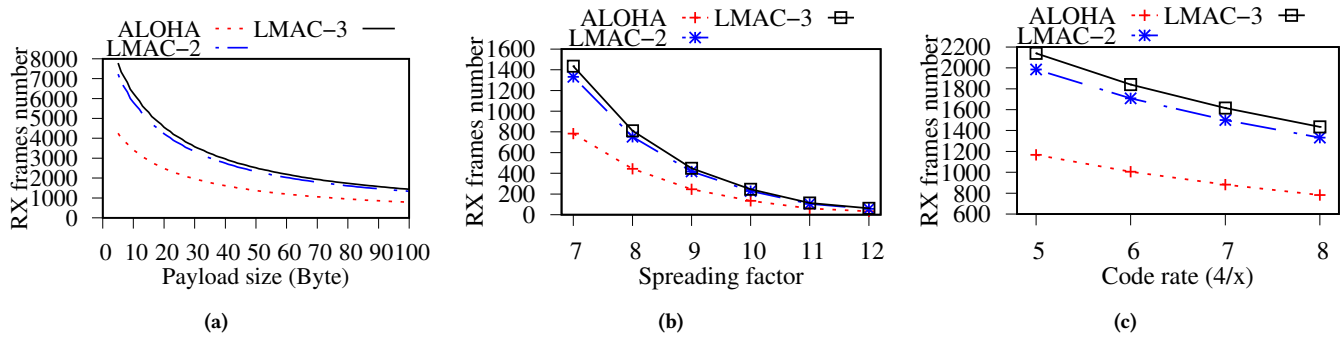


Figure 21: Number of successfully received frames in one minute vs. various parameters (a) varying payload size with SF7 and CR4/8; (b) varying SF with 16-byte payload and CR4/8; (c) varying CR with 16-bytes payload and SF7.

also show that LMAC-2 and LMAC-3 outperform ALOHA by 70% and 83% respectively, and the numbers of nodes that LMAC-2 and LMAC-3 could support would likely be similar.

8 RELATED WORKS

As the primitive ALOHA scales poorly with the number of LoRa nodes, several studies aimed at improving the resilience and throughput of LoRaWAN networks by physical (PHY) layer modification or new MAC designs. We now review these PHY and MAC layer techniques separately.

LoRa PHY. Choir [14] exploits the distinct frequency biases of end devices to disentangle the collided frames. As Choir uses a new PHY algorithm, the gateway of Choir has to use SDR, such as USRP, and employ software demodulation implementations. As such, Choir will lose the line speed of hardware gateways. Netscatter [19] reduces LoRa modulation to on-off keying to simplify the operation on backscatter devices. It can scale for concurrent demodulation of backscatter LoRa signals. However, the unique design of Netscatter’s on-off keying modulation does not conform to standard LoRa modulation. To address the significant attenuation of the signals transmitted from end devices deeply located within buildings, Charm [13] jointly processes the signals received by multiple gateways to increase received SNR and the network coverage. Although Charm does not require a full-fledged SDR device, dedicated programmable hardware design is needed to collect the I/Q data of the received signal. To address channel contention, Chime [16] analyzes a single frame from a node through multiple gateways. With gateway synchronization and optimal frequency estimation, a gateway advises the node on the optimal transmission channel. Chime requires SDR to act as gateway to collect the I/Q data and strict synchronization between gateways. Voigt et al. exploit the use of multiple LoRa gateways and directional antennae to aid in decoding interfered LoRa frames [33]. However, directional antennae impede a nodes’ ability to be heard by multiple gateways hence coverage of a node.

LoRa MAC. Slotted ALOHA and the more well scheduled time-division multiple access (TDMA) can reduce collisions compared with ALOHA. Specifically, the end node will transmit in certain time slots to alleviate the collision issue caused by random access in ALOHA. Several studies [18, 27, 32] focus on the clock synchronization service needed by slotted ALOHA and TDMA. However, due to LoRa’s limited communication bandwidth [17], the communications for clock synchronization present a burden for LoRaWAN

networks. Thus, CSMA is desirable in LoRaWANs, as it does not require clock synchronization. Beltramelli et al. present a stochastic geometry-based model [7] to analytically show that CSMA outperforms ALOHA in terms of reliability and energy efficiency. However, full-fledged implementation of CSMA in LoRaWANs is still unavailable due to the lack of hardware support for carrier sense. To achieve carrier sense, DeepSense [9] detects on-air transmission of LPWAN radios by storing spectrograms as an image, then performing signal identification on the image with artificial neural networks. However, an additional SDR device and machine learning accelerator are needed. To implement CSMA with off-the-shelf LoRa devices, a previous work [26] ports the 802.11 CSMA protocol to LoRa, in which the carrier sense is performed by a stand-alone node near the transmitter. This introduces additional monetary cost and system complexity. In this work, we leverage CAD for carrier sense on LoRa end devices and design LMAC to improve the goodput and PRR of LoRaWAN networks. LMAC requires no modifications to the LoRa PHY layer and the LoRaWAN specification. Thus, it can be readily deployed in the current LoRaWAN networks. To the best of our knowledge, this paper presents the first systematic research and implementation of CSMA for LoRa networks.

9 CONCLUSION

This paper shows that the Channel Activity Detection (CAD) feature of LoRa radios for preamble chirp detection can be exploited for reliable carrier sense. Based on CAD, this paper designs three versions of LMAC that respectively implements CSMA for LoRa networks, and balances the loads of the channels defined by frequencies and spreading factors by using the end nodes’ local information only and then additionally the gateway’s global information. Testbed experiments show that, compared with ALOHA, LMAC brings significant performance improvements in terms of PRR and goodput, as well as radio energy saving per successful frame delivery.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers and shepherd for their valuable comments and helpful suggestions that improved the quality of this paper. The work was supported by the Singapore MOE AcRF Tier 1 grant RG129/18, NTU CoE SUG grant 04INS000269C130, and Singapore MOE AcRF Tier 1 grant 2019-T1-001-044. The code of this work is built with the LMIC (LoRaMAC-in-C) library ported by Matthijs Kooijman.

REFERENCES

- [1] 2005. IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 15.1a: Wireless Medium Access Control (MAC) and Physical Layer (PHY) specifications for Wireless Personal Area Networks (WPAN). *IEEE Std 802.15.1-2005 (Revision of IEEE Std 802.15.1-2002)* (2005).
- [2] 2007. IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *IEEE Std 802.11-2007 (Revision of IEEE Std 802.11-1999)* (2007).
- [3] 2013. *SX1272/3/6/7/8: LoRa Modem*. Technical Report. Semtech Corporation. AN1200.13.
- [4] 2014. *Reading channel RSSI during a CAD*. Technical Report. Semtech Corporation. AN1200.21.
- [5] Ferran Adelantado, Xavier Vilajosana, Pere Tuset-Peiro, Borja Martinez, Joan Melia-Segui, and Thomas Watteyne. 2017. Understanding the limits of LoRaWAN. *IEEE Communications magazine* 55, 9 (2017).
- [6] Aloys Augustin, Jiazi Yi, Thomas Clausen, and William Mark Townsley. 2016. A study of LoRa: Long range & low power networks for the internet of things. *Sensors* 16, 9 (2016).
- [7] Luca Beltramelli, Aamir Mahmood, Patrik Österberg, and Mikael Gidlund. 2020. LoRa beyond ALOHA: An Investigation of Alternative Random Access Protocols. *arXiv preprint arXiv:2002.10732* (2020).
- [8] Martin C Bor, Utz Roedig, Thimo Voigt, and Juan M Alonso. 2016. Do LoRa low-power wide-area networks scale?. In *Proceedings of the International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWIM '16)*. ACM.
- [9] Justin Chan, Anran Wang, Arvind Krishnamurthy, and Shyamnath Gollakota. 2019. DeepSense: Enabling Carrier Sense in Low-Power Wide Area Networks Using Deep Learning. *arXiv preprint arXiv:1904.10607* (2019).
- [10] LoRa Alliance Technical Committee. 2017. *LoRaWAN™1.1 Specification*. Technical Report. LoRa Alliance.
- [11] 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 *Proceedings of the International Conference on Information Processing in Sensor Networks (IPSN '19)*. IEEE.
- [12] Jason A. Donenfeld. 2017. WireGuard: Next Generation Kernel Network Tunnel. (2017).
- [13] 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 *Proceedings of the International Conference on Information Processing in Sensor Networks (IPSN '18)*. IEEE.
- [14] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yaoundinedan. 2017. Empowering Low-Power Wide Area Networks in Urban Settings. In *Proceedings of Special Interest Group on Data Communication (SIGCOMM '17)*. ACM.
- [15] ETSI. 2018. *Short Range Devices (SRD) operating in the frequency range 25 MHz to 1 000 MHz; Part 2: Harmonised Standard for access to radio spectrum for non specific radio equipment*. Standard. STSI, Valbonne, FR.
- [16] Akshay Gadre, Revathy Narayanan, Anh Luong, Anthony Rowe, Bob Iannucci, and Swarun Kumar. 2020. Frequency Configuration for Low-Power Wide-Area Networks in a Heartbeat. In *Proceedings of the Symposium on Networked Systems Design and Implementation (NSDI '20)*. USENIX Association.
- [17] Branden Ghena, Joshua Adkins, Longfei Shangguan, Kyle Jamieson, Philip Levis, and Prabal Dutta. 2019. Challenge: Unlicensed LPWANs Are Not Yet the Path to Ubiquitous Connectivity. In *Proceedings of the International Conference on Mobile Computing and Networking (MobiCom '19)*. ACM.
- [18] Chaojie Gu, Rui Tan, and Xin Lou. 2019. One-Hop Out-of-Band Control Planes for Multi-Hop Wireless Sensor Networks. *ACM Transactions on Sensor Networks (TOSN)* 15, 4 (2019).
- [19] Mehrdad Hessar, Ali Najafi, and Shyamnath Gollakota. 2019. Netscatter: Enabling large-scale backscatter networks. In *Proceedings of the Symposium on Networked Systems Design and Implementation (NSDI '19)*. USENIX Association.
- [20] Jansen C. Liando, Amalinda Gamage, Agustinus W. Tengourtius, and Mo Li. 2019. Known and Unknown Facts of LoRa: Experiences from a Large-Scale Measurement Study. *ACM Transactions on Sensor Networks (TOSN)* 15, 2 (2019).
- [21] LMAC Integrated LoRaWAN. 2020. *arduino-lmic-csma*. <https://wands.sg/lmac>
- [22] Konstantin Mikhaylov, Juha Petaejaerä, and Tuomo Haenninen. 2016. Analysis of capacity and scalability of the LoRa low power wide area network technology. In *Proceedings of the European Wireless Conference (EW '16)*. VDE.
- [23] Konstantin Mikhaylov, Juha Petajajarvi, and Janne Janhunen. 2017. On LoRaWAN scalability: Empirical evaluation of susceptibility to inter-network interference. In *Proceedings of the European Conference on Networks and Communications (EuCNC '17)*. IEEE.
- [24] Sandeep Mistry. 2020. *Arduino-LoRa*. <https://github.com/sandeepmistry/arduino-LoRa>
- [25] Pierre Neumann, Julien Montavont, and Thomas Noel. 2016. Indoor deployment of low-power wide area networks (LPWAN): A LoRaWAN case study. In *Proceedings of the International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob '16)*. IEEE.
- [26] Congduc Pham. 2018. Robust CSMA for long-range LoRa transmissions with image sensing devices. In *Proceedings of Wireless Days (WD '18)*. IEEE.
- [27] Rajeev Piyare, Amy L. Murphy, Michele Magno, and Luca Benini. 2018. On-demand LoRa: Asynchronous TDMA for energy efficient and low latency communication in IoT. *Sensors* 18, 11 (2018).
- [28] Pieter Robyns, Peter Quax, Wim Lamotte, and William Thenaers. 2018. A multi-channel software decoder for the LoRa modulation scheme. In *Proceedings of the International Conference on Internet of Things, Big Data and Security (IoTBDs)*. SciTePress.
- [29] Moon Solutions. 2020. *High Voltage Power Monitor*. <https://www.msoon.com/online-store/High-Voltage-Power-Monitor-Part-Number-AAA10F-p90002590>
- [30] Jothi Prasanna Shanmuga Sundaram, Wan Du, and Zhiwei Zhao. 2019. A Survey on LoRa Networking: Research Problems, Current Solutions and Open Issues. *IEEE Communications Surveys & Tutorials* (2019).
- [31] Shuai Tong, Jiliang Wang, and Yunhao Liu. 2020. Combating Packet Collisions Using Non-Stationary Signal Scaling in LPWANs. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys '20)*. ACM.
- [32] Roman Trüb and Lothar Thiele. 2018. Increasing Throughput and Efficiency of LoRaWAN Class A. In *Proceedings of the International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM '18)*. IARIA.
- [33] Thimo Voigt, Martin Bor, Utz Roedig, and Juan Alonso. 2017. Mitigating inter-network interference in LoRa Networks. In *Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks (EWSN '17)*. Junction Publishing.
- [34] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2019. FTrack: parallel decoding for LoRa transmissions. In *Proceedings of the Embedded Networked Sensor Systems (SENSYS '19)*. ACM.
- [35] Cui Zhao, Zhenjiang Li, Ting Liu, Han Ding, Jinsong Han, Wei Xi, and Ruowei Gui. 2019. RF-Mehndi: A Fingertip Profiled RF Identifier. In *Proceedings of the International Conference on Computer Communications (INFOCOMM '19)*. IEEE.