Poster: Chirpbox – A Low-Cost LoRa Testbed Solution

Xiaoyuan Ma[1,2], Dan Li[1], Fengxu Yang[1,3], Carlo Alberto Boano[4], Pei Tian[1,2], Jianming Wei[1]
1. Shanghai Advanced Research Institute, Chinese Academy of Sciences, China
2. University of Chinese Academy of Sciences, China 3. ShanghaiTech University, China
4. Institute of Technical Informatics, Graz University of Technology, Austria
{maxy, lid, tianpei2018, wjm}@sari.ac.cn yangfx@shanghaitech.edu.cn cboano@tugraz.at

Abstract

Low-Power Wide Area Networks (LP-WANs) are becoming a key enabler of many practical IoT applications. Although LoRa is one of the most representative LP-WAN technologies, only a few large-scale LoRa testbeds are publicly open to network researchers, which limits the ability to test the real-world performance of LoRa-based solutions. To enable a large-scale testing on real hardware, in this poster we present Chirpbox, a low-cost LoRa testbed solution. Chirpbox aims to evaluate the different LoRa protocols rather than to debug a system in detail. By only using the LoRa radio transceiver on each testbed node, Chirpbox is able to: (i) send back the evaluation results, (ii) schedule (i.e., start and stop) experiments, as well as (iii) disseminate a firmware under test without the need of additional infrastructure, such as local area networks. Each node in Chirpbox can get synchronized with a GPS module and is able to compute time-related statistics such as end-to-end latency on its own. Moreover, in order to easily deploy a LoRa network over a large area, a battery set is integrated in each Chirpbox node, as there is no need for energy-hungry observers.

1 Motivation

LP-WANs are applied nowadays in many practical scenarios: from smart city and smart grid applications, to industrial monitoring and precise agriculture [3]. LoRa is one of the most representative technologies of LP-WANs. However, for network researchers, only a few open testbeds supporting LoRa are available to date. FIT-IoT [1] consists of 25 LoRa nodes within a 12 $m \times 12$ m indoor room. In FlockLab [5], there are only three LoRa nodes observed. Dandelion [6] has 51 LoRa nodes, but they cannot be accessed publicly so far. Except for these testbeds, it is hard to find another open facility supporting LoRa over a large area, ideally outdoors and with more than 20 nodes for large-scale experiments.

If nodes are deployed outdoors over a large area, in addition to the duty cycle limitations, a main technical challenge to test a LoRa system is the lack of infrastructure, i.e., a wired backbone that is commonly available in most indoor testbeds. Hence, it is hard for a tested LoRa-based system to send back the evaluation results without the availability of LAN/WLAN networks (**C1**); to disseminate a firmware under test (**C2**); to share a common time scale to compute time-related statistics such as latency (**C3**); to schedule an experiment remotely (**C4**); and to ensure long-term availability, as a stable power supply can not be guaranteed (**C5**). Moreover, the cost of a testbed node should ideally be minimized (**C6**).

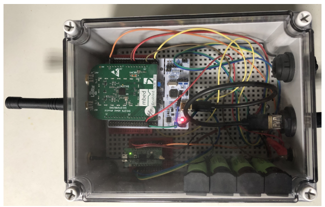

Figure 1. A node in Chirpbox.

To address these challenges, we propose Chirpbox, a low-cost LoRa testbed solution based on STM32 NUCLEO-L476RG with a SX1276MB1MAS LoRa module (Fig. 1). Unlike existing testbeds [1, 5, 6], there is no explicit energy-hungry observer for each node, which allows to address **C5**. Each node is equipped with four Li-ion rechargeable batteries (Nitecore NL1834). In order to solve **C4**, all the nodes are connected with a real-time clock (RTC) module[1] to control the duration of an experiment. Clock synchronization (**C3**) can be achieved by using a GPS module[2]. Since Chirpbox aims to evaluate different LoRa network protocols rather than to debug a system, the evaluation results do not contain much logging information. Therefore, after an experiment, each node sends back its evaluation result with the LoRa radio itself (**C1**). In Chirpbox, all the statistics are computed by nodes in a software manner, following the approach used by Energest in Contiki OS. Chirpbox uses a *control node* connected with a computer to schedule (i.e., start and stop) an experiment (**C4**) and to disseminate a firmware under test (**C2**) between two experiments. The control node is not actively used during experiments (i.e., it remains silent). For

---

[1] https://www.maximintegrated.com/cn/products/analog/real-time-clocks/DS3231MPMB1.html

[2] http://navspark.mybigcommerce.com/development-boards/

International Conference on Embedded Wireless Systems and Networks (EWSN) 2020
17–19 February, Lyon, France
ISBN: 978-0-9949886-4-5

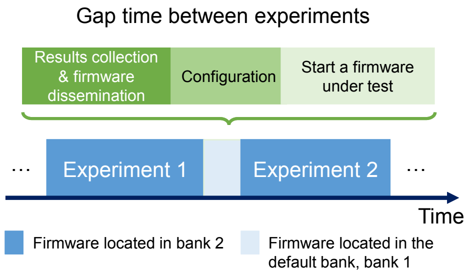

Figure 2. Chirpbox's workflow. A LoRa-based Mixer protocol is used to send back evaluation results to the control node and disseminate downlink information.

each node, the price of the main components (i.e., the GPS module, the RTC module, the LoRa node, and the battery set) is less than 160 US dollars (**C6**). Furthermore, no more Ethernet cables are required, which ensures a high deployment flexibility. In summary, Chirpbox can be easily deployed without any communication and power infrastructure, which drastically simplifies the testing of a LoRa-based system.

2 Highlights of Chirpbox

As mentioned in Sect. 1, there is no observer node in Chirpbox actively measuring and reporting data. Instead, two firmwares are stored on each node: the *daemon* and the firmware under test. In the daemon, we introduce a LoRa-based Mixer protocol (Sect. 2.1) to send evaluation results to the control node and to disseminate both scheduling information and firmware under test (results collection & firmware dissemination phase in Fig. 2). We use the dual-bank flash memory, i.e., two parallel flash blocks of the STM32L4 series MCU, to update and load the firmware under test on a node (configuration phase in Fig. 2). This feature ensures that writing in one bank will not interrupt reading (and fetching instructions) from the other bank. Moreover, the network topology can be obtained in the configuration phase.

2.1 Using a LoRa-based Mixer Protocol

Protocols based on concurrent transmissions (CT) can disseminate information at a very short latency. To exploit this property, we port Mixer [2], a CT-based all-to-all protocol, as well as an LBT & AFA (Listen Before Talk and Adaptive Frequency Agility) mechanism onto our LoRa nodes. The LoRa-based Mixer protocol is only called by the daemon between experiments. In the results collection & firmware dissemination phase (shown in Fig. 2), the results of an experiment are sent to the control node. Afterwards, in the configuration phase, the new firmware under test and the starting time of the next experiment are disseminated to all the nodes from the control node.

2.2 Updating a Firmware under Test

By using the dual-bank flash memory of the STM32L4 series MCU, we configure bank 1 as the default bank and locate the received new firmware in bank 2. The node runs the firmware located in the default bank after a hardware reset. Thus, the daemon embedding the LoRa-based Mixer is put in the default bank. The daemon runs during the gap time between two experiments as shown in Fig. 2. Specifically, the daemon writes the received firmware at the corresponding address of bank 2, configures the RTC module, and starts the firmware under test at the given time using the GPS module.

2.3 Patching a Firmware under Test

It is unnecessary to disseminate a new firmware if only a small portion differs from the previous one. Inspired by [4], we patch the firmware rather than disseminate the whole image. Specifically, we assume that all the tested nodes use one `.hex` file during each experiment and derive a `.diff` file describing the difference between the current binary file in bank 2 and the new firmware. The `.diff` file is then disseminated to all testbed nodes by the control node, so that they can modify the binary values at the specific addresses.

2.4 Topology and Health Status

Guaranteeing connectivity across the testbed is critical since both the collection of results and the dissemination of the firmware under test rely on LoRa communications. In the configuration shown in Fig. 2, Chirpbox is able to update the connectivity of the current topology with different spreading factors (SFs). When nodes update the topology, time is divided into slots (during which only one node transmits information) based on GPS information, and each slot is pre-assigned with a specific SF. This way, all nodes can obtain their connectivity tables and exchange them using the LoRa-based Mixer protocol. Please note that a topology update can be skipped if the LoRa network is known and static.

If the control node fails to collect evaluation results or the connectivity tables from all nodes, then malfunctioning is speculated by the control node and experiments are halted.

2.5 Redundant Nodes

Placing a few redundant nodes in the testbed that remain inactive during an experiment can be beneficial in several ways. On the one hand, these nodes allow to minimize the time between experiments. As redundant nodes did not take part in the experiments, they have less constraints on the number of transmissions they can carry out than the other nodes (who need to comply to the duty cycle limitations). Therefore, redundant nodes can be used to disseminate both the evaluation results and the firmware under test quicker through the network using the Lora-based Mixer protocol.

On the other hand, redundant nodes can also be used to monitor the LoRa signal and speculate the radio usage of the firmware under test. This way, any violation of LoRa's duty cycle limitations can be reported after an experiment.

3 References

[1] C. Adjih et al. FIT IoT-LAB: A Large Scale Open Experimental IoT testbed. In *Proc. of the 2nd WF-IoT World Forum*. IEEE Computer Society, December 2015.

[2] C. Herrmann, F. Mager, and M. Zimmerling. Mixer: Efficient Many-to-All Broadcast in Dynamic Wireless Mesh Networks. In *Proc. of the 17th SenSys Conference*. ACM, November 2018.

[3] U. Raza et al. Low Power Wide Area Networks: An Overview. *IEEE Communications Surveys and Tutorials*, 19(2), 2017.

[4] M. Schuß, C. A. Boano, and K. Römer. Moving Beyond Competitions: Extending D-Cube to Seamlessly Benchmark Low-Power Wireless Systems. In *Proc. of the 1st CPSBench Workshop*. IEEE, 2018.

[5] R. Trüb, R. D. Forno, T. Gsell, J. Beutel, and L. Thiele. A Testbed for Long-range LoRa Communication: Demo Abstract. In *Proc. of the 18th IPSN Conference*. ACM, April 2019.

[6] W. Wang, X. Guo, L. Liu, X. Lei, X. Zheng, M. Jin, and Y. He. Dandelion: Design of An Online Large Scale LoRa Testbed. In *Proc. of the 16th EWSN Conference*. Junction Publishing, Feburary 2019.