Poster: Towards a Federated Testbed Infrastructure for Geographically-Distributed Low-Power Wireless Systems

Markus Schuß*, Carlo Alberto Boano*, Michael Baddeley†, Monika Prakash†, and Kay Römer*
*Institute of Technical Informatics, Graz University of Technology, Austria; †Technology Innovation Institute TII, UAE
{markus.schuss, cboano, roemer}@tugraz.at, {michael.baddeley,monika.prakash}@tii.ae

Abstract

Several IoT testbeds have been designed to test and push the performance of low-power wireless networking protocols to the limit. However, they mostly target low-power wireless systems operating in isolation, and are unable to precisely characterize the performance of solutions operating across multiple sites or interacting with cloud resources. As advances in backbone communication networks allow to move towards decentralized IoT deployments, there is a growing need to understand the impact of the Internet on the timeliness and reliability of end-to-end communications. In this poster, we outline the necessary steps to evolve D-Cube, a full-fledged benchmarking infrastructure, into a *federated testbed* capable of measuring the performance of low-power wireless systems operating across multiple sites through the Internet, thereby enabling research on the next-generation IoT systems operating on a mesh-cloud continuum.

1 Introduction

The scalability offered by the cloud often plays a key role in today's IoT solutions: as a result, storage and computation are often hosted hundreds of kilometers away from the end-device. However, many IoT applications impose strict latency and reliability requirements (e.g., wireless control systems) on network performance, which can hardly be met when end-devices are distributed over large geographical areas and separated by large distances from the employed cloud resources. To cope with this, paradigms like edge and fog computing allow the creation of decentralized IoT solutions that forgo the benefits that are inherent to cloud solutions, but allow a higher responsiveness [6]. Such paradigms are often tied to the use of low-power wireless (LPW) protocols employing synchronous transmissions and complex duty-cycling strategies to enable a reliable and timely connection among end-devices over mesh networks [7].

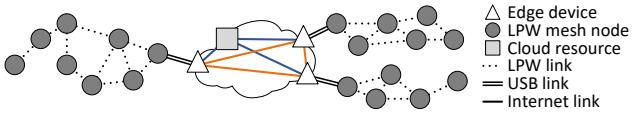

Figure 1. Example of a heterogeneous IoT architecture.

Fig. 1 shows an example of an IoT solution operating across multiple geographical sites. The links in blue show the use of the classical cloud paradigm, where the various sites are connected via a cloud resource (e.g., centralized MQTT). This may cause large delays, as the LPW nodes cannot directly communicate with their counterparts on another site. The addition of the links in orange tackles this problem, allowing edge devices to communicate directly while still retaining the possibility to leverage cloud resources.

In such a scenario, the ability to precisely test and characterize the end-to-end latency as well as reliability of communications is crucial. For example, several LPW protocols try to adjust their operations in order to maximize responsiveness, but their performance strongly depends, among others, on the congestion across the Internet and on the wireless technology used to interconnect the various sites [9].

Testbeds are commonly used to explore and evaluate the performance of IoT solutions. Testbeds tailored to LPW devices [4, 10], however, are mostly limited to a single location, and lack the ability to run code on edge devices / border routers (BRs) and to facilitate connectivity to other sites. Federated testbeds [2, 5] unify the access to multiple testbeds via a common API and have even been set up across multiple continents [1], but lack the means to collect measurements in hardware (e.g., energy consumption) from the LPW devices. A notable exception is FIT IoT-LAB [3], which enables hardware measurements, but lacks public IP addresses or the ability to forward ports from the Internet to the edge devices, which are needed for incoming connections.

Open challenges. To characterize IoT solutions operating across multiple sites, it is necessary to go beyond traditional federated testbeds and tackle the following problems.

Lack of automation. The APIs of federated testbeds allow control of the devices at each site, but there is no automated approach enabling the execution of an experiment *across* sites. All steps – from the setup of all devices (e.g., BRs on edge devices), the configuration of the network topology (e.g., addresses and routes), and the collection of measurements, to the synchronized start of the experiment – have to be manually performed by the user.

Lack of an isolated network. Opening up the edge devices to the Internet puts them at risk of attacks, and undesired traffic may be forwarded to the LPW network. Moreover, many protocols for LPW devices rely on IPv6 addresses that may not be available, requiring the use of address translation.

Lack of security. Federated testbeds offer access to edge devices (via virtualization) and support arbitrary code execution, but allowing users to run software on machines connected to the Internet – especially via public IP addresses – has the potential for abuse. A user could use the machines to share illicit material or execute denial of service attacks.

2 Enabling Testbed Federation

We aim to federate two testbeds based on D-Cube: the existing public instance in Graz, AT (`https://iti-testbed.tugraz.at`), and a private instance in Abu Dhabi, UAE (not accessible from the Internet), with both instances operating independently when no federated experiment is conducted. D-Cube already supports GPS timestamped measurements and automatically collects key performance metrics [8] using its Linux-based *observer modules*; currently, however, the user only has access to the LPW devices via an automated workflow. In addition to adding a second instance, federation requires key changes to D-Cube's architecture:

(i): Each instance's scheduler operates independently and the exact time at which an experiment runs depends on the number of other users in the queue. The scheduler needs to be extended to allow coordination of experiments across testbeds.

(ii): An overlay network linking the BRs at all sites needs to be automatically set up for each experiment. This allows LPW devices to communicate with devices in other testbeds without being interfered by foreign traffic.

(iii): Running arbitrary code compromises the security of the testbed: as such, user-provided software must be run in an isolated environment on the observer module, and given only limited permissions (e.g., to access the LPW node).

Containerized services. As there is no common BR, each solution has its own approach running on the edge device and forwarding data from/to the LPW network. Running the BR software directly on the observer module may compromise the testbed, as the software could alter the operating system (OS) or access the testbed's control network. It is hence necessary to run the user's code inside an isolated environment. Using virtualization to run a separate OS, however, would introduce delays and require a large amount of dedicated resources. A suitable alternative are Linux containers, which allow bundling all the necessary software into an image.

Overlay network. Exploring the impact of the Internet on mesh-cloud-mesh communication (i.e., relaying data from one mesh to another), while isolating the experiment from unwanted traffic can be achieved via a virtual private network (VPN). However, careful consideration needs to be put into the selection of the VPN solution, as it may introduce additional latency and alter the lossy nature of Internet links by encapsulating traffic in reliable TCP packets. A layer 2 VPN can be used to allow all BRs to directly communicate with one another. This also gives the user full control over addresses and routes without compromising the existing network's security. An observer module hence connects the BR container exclusively to this overlay network and leaves the configuration of the network interface to the container.

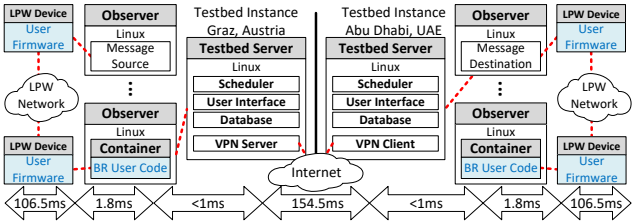

Figure 2. Two federated D-Cube instances running a joint experiment. The items in blue are provided by the user. The data flow (source to destination) is marked in red.

Federated scheduler. In addition to selecting the desired LPW platform from the attached devices and flashing the user-provided firmware, D-Cube also has to set up the containers with the user's image on the observer modules acting as BRs, and connect them with the overlay network isolating the experiment's traffic. These modifications allow D-Cube to autonomously benchmark the performance of solutions utilizing the Internet for a vital part of their communication.

3 Preliminary Results & Next Steps

Fig. 2 shows the two D-Cube instances running a joint experiment in which the source and destination of messages are located at different sites. We break down the latency of the individual segments when using Contiki-NG running on D-Cube's nRF52840DK targets. We measure an average latency across the two continents of 154.5±12.1 ms. The Ethernet connection between observer and server adds <1 ms. The nature of the USB connection between LPW device and BR plays a much larger role: the USB-to-UART converter adds 18 ms latency versus 1.8 ms of native USB. The LPW devices run 6TSCH, which adds 106.5±55.3 ms for each hop.

Future work. Our initial evaluation shows that our proposed approach is feasible. Still, some of the steps outlined in § 2 (e.g., the configuration of the network) are not yet fully automated. We plan next to finalize the implementation and to characterize the end-to-end performance of decentralized IoT solutions between the two sites in an automated manner.

Acknowledgements. This work was performed within the SPiDR project and was partially supported by the TU Graz LEAD project "Dependable IoT in Adverse Environments".

4 References

[1] EU-BR FUTEBOL project. `http://futebol.inf.ufrgs.br`.
[2] Fed4FIRE+. `https://www.fed4fire.eu`.
[3] C. Adjih et al. FIT IoT-LAB: A Large-Scale Open Experimental IoT Testbed. In *Proc. of the 2nd WF-IoT Forum*, 2015.
[4] P. Appavoo et al. Indriya2: A Heterogeneous Wireless Sensor Network (WSN) Testbed. In *Proc. of the 13th TridentCom Conf.*, 2018.
[5] Crew project. Testbeds, 2010. `http://www.crew-project.eu`.
[6] M. De Donno et al. Foundations and Evolution of Modern Computing Paradigms: Cloud, IoT, Edge, and Fog. *IEEE Access*, 7, 2019.
[7] M. Schuß et al. A Competition to Push the Dependability of Low-Power Wireless Protocols to the Edge. In *Proc. of EWSN'17*.
[8] M. Schuß et al. Moving Beyond Competitions: Extending D-Cube to Seamlessly Benchmark LPW Systems. In *Proc. of CPSBench'18*.
[9] M. Spörk et al. Ensuring End-to-End Dependability Requirements in Cloud-based Bluetooth Low Energy Appl. In *Proc. of EWSN'21*.
[10] R. Trüb et al. FlockLab 2: Multi-modal Testing and Validation for Wireless IoT. In *Proc. of the 3rd CPS-IoTBench Workshop*, 2020.