

Poster: Making D-Cube an Open Low-Power Wireless Networking Benchmark

Markus Schuß, Carlo Alberto Boano, and Kay Römer

Institute of Technical Informatics, Graz University of Technology, Austria
{markus.schuss, cboano, roemer}@tugraz.at

Abstract

To facilitate a reproducible benchmarking of low-power wireless systems, open-source tools and an agreed-upon set of scenarios and metrics are necessary. To this end, in 2016, we created D-Cube, a low-cost measurement tool that can accurately measure in hardware key performance metrics such as end-to-end delay, reliability, and power consumption. Through the past four years, D-Cube has mainly been used to support a dependability competition comparing the performance of state-of-the-art protocols under the same settings. In this poster, we outline the recent evolution of D-Cube from a distributed measurement tool used in public benchmarking events to an *open benchmark* available to the community as a common yardstick for evaluating and comparing the performance of low-power wireless systems.

1 Introduction

As an increasing number of Internet of Things systems imposing strict dependability requirements on network performance is being developed and commercialized, the demand for dependable low-power wireless protocols delivering data in a reliable, efficient, and timely manner is raising. In response to this need, the research and industrial community has proposed several solutions over the last decade. However, the lack of a standardized methodology to evaluate protocol performance [3] and the limited ability to systematically test low-power wireless protocols under the same settings often leads to a high divergence across experimental setups, which makes it hard, if not impossible, to compare different solutions. This raises the need for tools enabling a rigorous benchmarking of low-power wireless systems [2].

To bridge this gap, we have developed D-Cube, a low-cost measurement tool that can augment existing low-power wireless testbeds by accurately measuring in hardware key performance metrics such as end-to-end delay, reliability, and

power consumption [5]. To evaluate and compare the performance of state-of-the-art protocols, we started in 2016 a dependability competition co-located with the EWSN conference [1]. During the competition, we examined the dependability of different solutions by setting up a testbed infrastructure based on D-Cube, and awarded those sustaining the best performance. We replicated this event in the following three years; each time with a different evaluation scenario, which varied from point-to-point communication of single events to multi-byte data collection and dissemination.

While this format was well-received and helped developing low-power wireless protocols highly-resilient to RF noise, we received feedback from the community about a few shortcomings resulting from the competition's format:

Closed access. The competition's infrastructure based on D-Cube was only available to the contestants during a short time window (i.e., the competition's preparation phase), and could not be used by anyone else beyond this time span.

Single outdated platform. The competition was based on TelosB nodes, a platform created in 2004 that is widespread, supported in most public testbeds, and for which an implementation of several Glossy- and TSCH-based protocols exists. However, with the rise of more capable platforms supporting multiple standards, it becomes important to benchmark low-power wireless protocols on newer platforms, ideally co-located to older ones for a direct comparison.

No long-term comparisons. The competition's evaluation scenario changed across years. These changes were reflected also in the testbed and hardware configuration, which made a comparison between the performance of a new protocol and the previous competition's results impossible.

Single instance of D-Cube. Replicating D-Cube at other institutions would be beneficial to increase the set of scenarios available to the community. Unfortunately, whilst D-Cube's hardware design is open (<http://www.iti.tugraz.at/D-Cube>), the same does not hold true for its software.

To mitigate these issues, we have worked towards making D-Cube an open benchmark, and adapted its hardware and software accordingly, as described in the next section.

2 Towards an Open Benchmark

To make D-Cube an open benchmark and ensure its long-term availability, several changes to D-Cube were necessary:

(i) To allow long-term comparisons, we implemented a *permanent benchmarking* feature, i.e., the ability to create a




Figure 1. D-Cube observer node with two target nodes attached via a custom multiplexer (Mux) board.

database of protocols that are permanently under test. This way, protocols are re-run periodically to account for possible re-location of nodes, changes in office occupancy, or any other factor that could skew the performance comparisons.

(ii) In the past, the testbed was reconfigured yearly and only one evaluation scenario was available at a given time. To overcome this limitation, we enabled the creation of several *benchmark suites*, i.e., specific application scenarios targeting a certain hardware platform. Users can now evaluate a protocol on a given benchmark suite at any time.

(iii) The original D-Cube design supported only one target platform per observer node. We hence manufactured a *multiplexer* board such that multiple target platforms can be connected to an observer node at the same time.

(iv) To simplify the replication of D-Cube at other institutions without having to care about software dependencies or network configuration, all D-Cube components are now managed through Docker (<https://www.docker.com/>).

Fig. 2 outlines the workflow of D-Cube, which is an updated version of our previous efforts [4]. Developers can create a protocol for a specific *benchmark suite* using the Web interface and upload a firmware in hex format with a placeholder for the experiments’ parameters at a well-known location. The benchmark suite contains information on the target node, as well as the available settings for traffic load and generated RF-interference. The testbed automatically injects the values of these parameters into the firmware using binary patching and then executes it. The collected measurements are interpreted based on the benchmark suite’s description of output metrics. Once finalized, the protocol will be tested *perpetually* to ensure long-term comparability of results.

Changes to D-Cube’s hardware. D-Cube is built around a low-cost custom PCB and an off-the-shelf Raspberry Pi 3B, forming an observer node as shown in Fig. 1. Any target node that can be powered and programmed via USB (e.g., a TelosB, TI CC2650 Launchpad, OpenMote B or Nordic’s NRF52840-DK) can be easily benchmarked simply by being attached to an observer. To facilitate the comparison between different platforms, D-Cube was extended by a low-cost *multiplexer* board and now supports up to four nodes. This multiplexer is built using general purpose analog switches for the GPIO pins as well as specialised variants for I²C and USB. In addition, power switches are used to turn off all but one target node at a time. The observer node can now




Figure 2. D-Cube’s benchmarking workflow.

select a single target node to be powered and instrumented without any influence from other nodes attached to it.

Changes to D-Cube’s software. To ease D-Cube’s replication, the software has been built on top of Docker. The latter enables others to deploy D-Cube without having to care about software dependencies or network configuration, as Docker takes care of the communication between components. Only the ports for measurement data input, communication with the observer modules and the user-facing web-interface are exposed. The messaging between server and observer nodes is implemented using the Advanced Message Queuing Protocol, which allows us to implement a remote procedure call architecture keeping the workflow on the server-side agnostic to the programming of the target nodes.

3 Conclusions

In this poster, we have outlined the recent evolution of D-Cube from a distributed measurement tool used in public benchmarking events to an open benchmark available to the community as a common yardstick for evaluating and comparing the performance of low-power wireless systems. D-Cube is available at <https://iti-testbed.tugraz.at>.

Acknowledgments

This work has been performed within the LEAD project “Dependable Internet of Things in Adverse Environments” funded by Graz University of Technology and has been partially supported by the FFG-funded Pro²Future COMET center (contract no. 6112792, <https://pro2future.at/>).

4 References

- [1] C. A. Boano et al. EWSN Dependability Competition: Experiences and Lessons Learned. *IEEE Internet of Things Newsletter*, 2017.
- [2] C. A. Boano et al. Towards a Benchmark for Low-power Wireless Networking. In *Proc. of the 1st CPSBench Workshop*. IEEE, 2018.
- [3] R. Jacob, C. Boano, U. Raza, M. Zimmerling, and L. Thiele. Towards a Methodology for Experimental Evaluation in Low-Power Wireless Networking. In *Proc. of the 2nd CPS-IoTBench Workshop*. ACM, 2019.
- [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] M. Schuß et al. A Competition to Push the Dependability of Low-Power Wireless Protocols to the Edge. In *Proc. of the 14th EWSN Conf.*, 2017.