Dataset: Wireless Link Quality Estimation on FlockLab – and Beyond

Romain Jacob*
ETH Zurich
Switzerland
jacobr@ethz.ch

Reto Da Forno
ETH Zurich
Switzerland
rdaforno@ethz.ch

Roman Trüb
ETH Zurich
Switzerland
rtrueb@ethz.ch

Andreas Biri
ETH Zurich
Switzerland
abiri@ethz.ch

Lothar Thiele
ETH Zurich
Switzerland
thiele@ethz.ch

ABSTRACT

In wireless networks, link quality quantifies how "challenging" the environment is for a communication protocol. The performance of protocols are often compared using wireless testbeds and it is therefore important to monitor the testbed's link quality to ensure a fair comparison ground across different protocols.

Thus, we are collecting link quality estimation data for FlockLab, an extensively used public testbed. We publish this dataset together with our data collection firmware. This firmware has been designed to facilitate the collection of similar datasets for other wireless networks. The dataset and firmware are publicly available together on Zenodo: doi.org/10.5281/zenodo.3354718.

CCS CONCEPTS

• **Networks → Network measurement**.

KEYWORDS

dataset, wireless link quality, long-term monitoring

ACM Reference Format:
Romain Jacob, Reto Da Forno, Roman Trüb, Andreas Biri, and Lothar Thiele. 2019. Dataset: Wireless Link Quality Estimation on FlockLab – and Beyond. In *The 2nd Workshop on Data Acquisition To Analysis (DATA'19), November 10, 2019, New York, NY, USA*. ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3359427.3361907

1 INTRODUCTION

Link quality refers to how "good" the physical connection between two nodes is. The definition of "good" depends on the context: for wireless links, quality is often measured with a reliability metric.

Generally, wireless links are volatile, which makes it difficult to confidently estimate their quality. Links may be degraded (or even completely suppressed) by sporadic interference from other networks. As interference is typically unpredictable and unavoidable, one may see wireless link quality as a probabilistic process.

Knowing about link quality is important when evaluating and comparing the performance of communication protocols. A sound performance evaluation requires experiments to be *reproducible*. Formalizing reproducibility in networking is non-trivial (see *e.g.*, [6]) but one thing is clear: repeatability requires that experimental conditions (including link quality) are comparable between tests.

The best way to validate that link quality remains similar would be to *measure it during experiments*. Unfortunately, this is not possible: the communication protocol under test itself occupies the links, making it difficult to isolate the sole impact of external interference.

Another approach is to *profile the network* to learn the distribution of the link quality over time. In particular, the distribution may exhibit some patterns, called *seasonal components*. As wireless interference is due to cross-traffic from other networks, the link quality is expected to be "better" when there is less traffic. For example, in an office building, there is more cross-traffic generated during office hours; link quality is then expected to be "worse". Thus, daily and weekly seasonal components are expected.

To compare protocol performance, one must know and account for such seasonal components. If seasonal effects are neglected, experiments may run in different conditions (*e.g.*, with consistently more or less interference). Then, any observed performance difference may be only an artifact from testing under different conditions, and is therefore inconclusive.

Researchers rely on testbeds to develop, test, and compare wireless protocols; yet, the link quality distributions are rarely investigated. FlockLab [8] is one such testbed; it is an indoor network comprising 27 nodes spread across an entire office floor.[1] [2]

We aim to enable repeatable and sound performance comparison on FlockLab. Thus, we collect data to quantify link quality and identify potential seasonal components. As FlockLab is located in an office building, we expect both daily and weekly patterns (correlated with office hours). Detecting daily patterns requires a time resolution of measurements on an hourly scale. The dataset we present in this abstract contains results of link quality measurements performed on FlockLab *every two hours*. To date, data have been collected for *two months* and collection is continuing.

*Corresponding author; CRediT for authors' contributions available online [7].

[1] A snapshot of the FlockLab floor plan is available at https://tiny.cc/flocklab.
[2] Similar testbeds include *e.g.*, D-Cube [9], FIT IoT-LAB [2] and Twonet [11].

DATA'19, November 10, 2019, New York, NY, USA

ACM ISBN 978-1-4503-6993-0/19/11…$15.00
https://doi.org/10.1145/3359427.3361907

In addition, we publish the firmware used for collecting this dataset. All networks are different and link quality distributions depend on location-specific interference conditions, which are difficult to estimate. Therefore, to compare protocols on other testbeds than FlockLab, one should go through the same process of collecting link quality data to profile the environment. Thus, we design our data collection firmware to facilitate the collection of similar datasets for other wireless networks.

The rest of this abstract presents the dataset (Sec. 2) and the collection firmware (Sec. 3); both of which are publicly available [7].

2 THE DATASET

We collect wireless link quality data for FlockLab [8], an indoor network testbed comprising 27 nodes spread across an entire floor.[1] This section first briefly discusses the rationale for this dataset, describes the dataset itself, and finally presents some insights we gained from it. The dataset is publicly accessible on Zenodo [7].

2.1 Looking for Seasonality

In ongoing work, we are interested in the time dependencies of link quality in wireless networks, in particular seasonal components and long-term trends. As discussed in the introduction, knowing these dependencies is key for reproducible networking experiments [6].

The FlockLab testbed [8] has been running since 2012. Over the past seven years, link quality data have been collected, but they lack the regularity and temporal resolution necessary to investigate short-term seasonality such as daily and weekly patterns. This motivated the collection of a new dataset, described below.

2.2 Scenario

The data collection scenario is simple: Each node is assigned one dedicated time slot. In this slot, a node sends 100 packets, called *strobes*. All strobes have the same payload size and use a given radio frequency channel and transmit power.[3] All other nodes listen for the strobes and log packet reception events (*i.e.*, *success* or *failed*).[4] The assignment of nodes to slots is pseudo-random to avoid introducing correlation between tests.

The scenario runtime is short (tens of seconds). On FlockLab, the entire scenario takes 65 s for 27 nodes sending an 8-bytes payload. In fact, most of the runtime is used for network time synchronization.[5] Sending the strobes only takes a few seconds.

2.3 Dataset

The test scenario (see Sec. 2.2) is being run every two hours on the FlockLab testbed [8], for both TelosB [3] and DPP-cc430 [4] platforms. At the time of writing, we collected data from $\approx$ 500 tests per platform spread over two months, using all 27 nodes currently available.[1] We use channels 26 (2.48 GHz) and 5 (869 MHz) for the TelosB and DPP-cc430, respectively.

The dataset is available online [7] and contains:

- All FlockLab results, including serial logs and GPIO traces,
- The pre-processed data, stored in `.csv` files (one file per platform and per month, as described below),
- The script used to pre-process the data,
- The data collection firmware (see Sec. 3 for more details),
- The required files and information to patch the firmware,
- Data visualization plots, including those used in this abstract,
- This abstract.

The pre-processed `.csv` files structure the data as follows:

date_time	Test date and time in UTC format
test_number	FlockLab test number
rf_channel	Radio frequency channel (*i.e.*, 11 to 26 for the 2.4 GHz band; 0 to 10 for the sub- GHz band),
tx_power	Transmission power in dBm,
payload	Strobe payload size in bytes
host_id	ID of the node used for time synchronization
rand_seed	Seed used for randomizing the assignment of nodes to slots
snd_id	ID of the sending node (*i.e.*, the one strobing)
rcv_id	ID of the receiving node
rcv_total	Number of strobes from **snd_id** successfully received by **rcv_id**
rcv_stream	Bit-stream of packet reception events from **snd_id** by **rcv_id** (*e.g.*, '1101...' indicates that the third strobe has been lost)

Thus, for each test, there are two rows of data per pair of nodes; one row per direction.

2.4 Dataset Updates

The collected data already provide some insights on the short-term seasonality (see Sec. 2.5). Investigating long-term effects requires to continue the data collection, which we intend to do for the foreseeable future.[6] We envision monthly updates of the repository [7].

2.5 Preliminary Dataset Analysis

Link quality data allow to easily compute a *connectivity matrix* of the network, *i.e.*, a quantification of the (estimated) link quality between pairs of nodes. Fig. 1 shows an example matrix. It is produced by the pre-processing script included with the dataset [7].

Furthermore, thanks to the high temporal resolution of the dataset, one can investigate short-term seasonal components in link quality on FlockLab. The hypothesis is that there is more wireless interference during office hours, which should result in daily and weekly patterns in the data.

The expected daily and weekly seasonal components are indeed clearly visible in the TelosB data (Fig. 2a). However, there is no apparent seasonality for the DPP-cc430 network (Fig. 2b). The difference between the two can be easily explained: the TelosB operates in the 2.4 GHz band, where there is interference from WiFi and Bluetooth devices (thus correlated with office hours). Conversely, the DPP-cc430 operates on the sub- GHz band; in this frequency band, there is also interference with other technology (*e.g.*, LoRa) but it appears to be less time correlated; at least nowadays.

[3]These are tunable parameters, see Sec. 3.

[4]No other data (*e.g.*, RSSI, SNR) are collected. We argue that packet reception is the most direct and most reliable metric for wireless link quality at the network layer. Furthermore, collecting only packet reception events makes the firmware immediately portable to other platforms (thanks to *Baloo*) whereas *e.g.*, RSSI measurements are inherently radio-specific.

[5]The need for time synchronization is discussed in Sec. 3.

[6]However, we will likely decrease the time resolution to free up testing time.

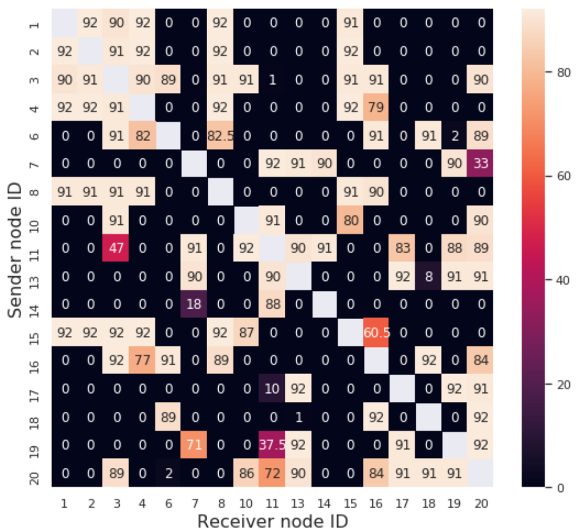

Figure 1: (Partial) FlockLab connectivity matrix for the TelosB platform. *The matrix shows the median number of strobes received between each pair of nodes in tests run in July 2019. For visibility, we limit the visualization to node 1 to 20.* The figure is "clickable" and available in the dataset repository [7].

3 THE DATA COLLECTION FIRMWARE

We are happy to share the link quality data we collect for FlockLab, but we also want to make it easier for others to collect similar datasets for other wireless networks. To achieve this, we also release our collection firmware, designed to be easily reusable.

Our approach consists in controlling the data collection entirely in software. The idea is to avoid relying on testbed infrastructure (*e.g.*, external triggers from FlockLab) such that the resulting firmware can be used on any wireless network.

The software is implemented using *Baloo* [5], a flexible network stack design framework based on synchronous transmissions. *Baloo* makes the programming easy; the framework handles timers and time synchronization, it provides a rich yet simple programming interface to implement the logic of the communication protocol, and it features out-of-the-box portability: a single protocol implementation lets the user compile firmware for all platforms supported by the framework. See [1] for more details.

To reuse our firmware on other networks (and/or for other purposes), the user can simply patch a set of experiment parameters directly in the firmware; *i.e.*, parameters can be modified without recompiling the source code. Currently, the following parameters can be patched: **rf_channel**, **payload**, **host_id**, and **rand_seed** (see Sec. 2.3 for the parameter description).

The firmware source code is available in the *Baloo* repository [1]. Compiled firmware for both TelosB [3] and DPP-cc430 [4] platforms are available together with the dataset [7]. Code and firmware for the nRF52840 [10] is under development and will be released soon.

[7] We consider a link as "good" when at least 50 strobes are received. This definition is arbitrary: one may choose many different metrics for link quality which would capture various properties of interest (*e.g.*, the burstiness of links). We make this possible by including the raw data in the dataset.

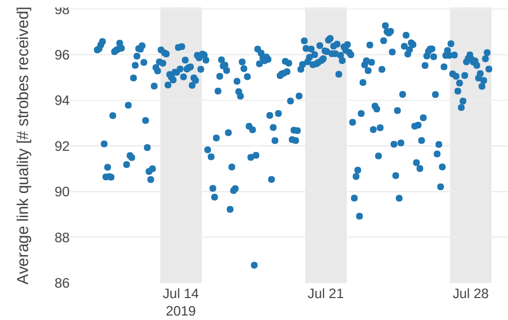

(a) TelosB. *Channel* 26 *(2.48 GHz),* 8*-bytes payload,* 0 *dBm TX power.*

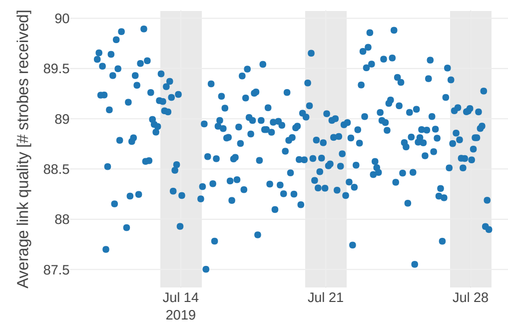

(b) DPP-cc430. *Channel* 5 *(869 MHz),* 8*-bytes payload,* 0 *dBm TX power.*

Figure 2: Time series of link quality on FlockLab for the TelosB (2a) and DPP-cc430 (2b) platforms. Data points show the mean number of strobes received on all "good" links.[7] Shaded areas correspond to weekends. *Fig. 2a clearly shows the expected weekly and daily seasonal components: the link quality consistently drops during day time, except on the weekends. Conversely, this pattern is not visible on Fig. 2b. The difference can be explained by the different frequency bands used by the TelosB and DPP-cc430 (2.4 GHz and sub- GHz respectively). The data show that wireless interference is more correlated with human activity in the 2.4 GHz band, due* e.g.*, to WiFi and Bluetooth devices.* The figures are "clickable" and available in the dataset repository [7].

4 CONCLUSION

Guaranteeing reproducibility is fundamental in science, but appears very challenging to achieve in practice. Wireless networking is no exception. Arguably, the situation may even be particularly difficult, due to the unpredictable and hardly controllable variability of wireless links.

Investigating and accounting for link variability is necessary to enable fair comparisons of wireless protocols performance. Therefore, we collect and publish a dataset enabling link quality estimation for the FlockLab testbed; in particular we focus on identifying seasonal effects (*i.e.*, periodic patterns). Furthermore, we design a data collection firmware that lets other users easily collect similar datasets for other wireless networks.

We release our dataset and firmware with the hope of facilitating the monitoring of experimental conditions in wireless networking, and henceforth contributing to improve reproducibility in the field.

REFERENCES

[1] 2018. Baloo. http://www.romainjacob.net/baloo/. [Online] - Last accessed: 2019-10-11.

[2] C. Adjih, E. Baccelli, E. Fleury, G. Harter, N. Mitton, T. Noel, R. Pissard-Gibollet, F. Saint-Marcel, G. Schreiner, J. Vandaele, and T. Watteyne. 2015. FIT IoT-LAB: A Large Scale Open Experimental IoT Testbed. In *2015 IEEE 2nd World Forum on Internet of Things (WF-IoT)*. 459–464. https://doi.org/10.1109/WF-IoT.2015.7389098

[3] Advanticsys. [n. d.]. MTM-CM5000-MSP 802.15.4 TelosB Mote Module. https://www.advanticsys.com/shop/mtmcm5000msp-p-14.html. [Online] - Last accessed: 2019-10-11.

[4] Texas Instruments. [n. d.]. CC430F6137 16-Bit Ultra-Low-Power MCU. http://www.ti.com/product/CC430F6137. [Online] - Last accessed: 2019-10-11.

[5] Romain Jacob, Jonas Bächli, Reto Da Forno, and Lothar Thiele. 2019. Synchronous Transmissions Made Easy: Design Your Network Stack with Baloo. In *Proceedings of the 2019 International Conference on Embedded Wireless Systems and Networks*.

[6] Romain Jacob, Carlo Alberto Boano, Usman Raza, Marco Zimmerling, and Lothar Thiele. 2019. Towards a Methodology for Experimental Evaluation in Low-Power Wireless Networking. In *2nd Workshop on Benchmarking Cyber-Physical Systems and Internet of Things (CPS-IoTBench'19)*. https://doi.org/10.3929/ethz-b-000325096

[7] Romain Jacob, Reto Da Forno, Roman Trüb, Andreas Biri, and Lothar Thiele. 2019. Wireless Link Quality Estimation on FlockLab - and Beyond. https://doi.org/10.5281/zenodo.3354718

[8] Roman Lim, Federico Ferrari, Marco Zimmerling, Christoph Walser, Philipp Sommer, and Jan Beutel. 2013. FlockLab: A Testbed for Distributed, Synchronized Tracing and Profiling of Wireless Embedded Systems. In *Proceedings of the 12th International Conference on Information Processing in Sensor Networks (IPSN '13)*. ACM, New York, NY, USA, 153–166. https://doi.org/10.1145/2461381.2461402

[9] Markus Schuß, Carlo Alberto Boano, and Kay Römer. 2018. Moving Beyond Competitions: Extending D-Cube to Seamlessly Benchmark Low-Power Wireless Systems. In *CPSBench 2018 - International Workshop on Benchmarking Cyber-Physical Networks and Systems*. 6.

[10] Nordic Semiconductors. [n. d.]. nRF52840. https://www.nordicsemi.com/Products/Low-power-short-range-wireless/nRF52840. [Online] - Last accessed: 2019-10-11.

[11] S. Yin, O. Gnawali, P. Sommer, and B. Kusy. 2014. Multi Channel Performance of Dual Band Low Power Wireless Network. In *2014 IEEE 11th International Conference on Mobile Ad Hoc and Sensor Systems*. 345–353. https://doi.org/10.1109/MASS.2014.120