

# The Impact of Decreasing Transmit Power Levels on FlockLab To Achieve a Sparse Network

Matthew Bradbury  
WMG  
University of Warwick  
M.Bradbury@warwick.ac.uk

Arshad Jhumka  
Department of Computer Science  
University of Warwick  
H.A.Jhumka@warwick.ac.uk

Carsten Maple  
WMG  
University of Warwick  
CM@warwick.ac.uk

## ABSTRACT

For research in Wireless Sensor Networks (WSNs) and the Internet of Things (IoT), while many protocols are either analysed mathematically or simulated to assess performance, it has become necessary that they are tested on hardware in real world environments. Algorithms are often validated in either (i) a densely connected network or (ii) a sparsely connected network. The majority of existing testbeds have implemented dense networks, making evaluation and validation of certain protocols, such as spatially-redundant source location privacy (SLP) protocols, challenging. We explore the use of transmission power to achieve network sparsity and present the results of such experiments performed on the FlockLab testbed. Based upon our experience in using FlockLab, we also identify additional aspects of testbed performance that need to be monitored to ensure reliable and reproducible results.

## CCS CONCEPTS

• **Networks** → **Network measurement; Sensor networks.**

## KEYWORDS

Benchmark, Wireless Sensor Networks, IoT, FlockLab

### ACM Reference Format:

Matthew Bradbury, Arshad Jhumka, and Carsten Maple. 2019. The Impact of Decreasing Transmit Power Levels on FlockLab To Achieve a Sparse Network. In *2nd Workshop on Benchmarking Cyber-Physical Systems and Internet of Things (CPS-IoTBench '19)*, April 15, 2019, Montreal, QC, Canada. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3312480.3313171>

## 1 INTRODUCTION

Research in Wireless Sensor Networks and the Internet of Things has reached the point where it is no longer sufficient to only simulate a new algorithm to demonstrate its effectiveness. The new technique must be deployed on real world hardware where the results demonstrate its actual effectiveness. With new requirements on novel applications which are expected to be deployed over large outdoors areas, testbeds may need to have specific properties, such as network sparsity and being located outdoors. However, an issue with this is that there are a limited number of testbeds available

(including FIT IoT-LAB [1], FlockLab [10], Indriya2 [2]). Further, these testbeds are located indoors with a node layout that cannot be customised by their users. There are some testbeds based outdoors, such as AWSP [14], but it is unclear how other researchers can deploy experiments to this testbed. A further outside WSN testbed, PotatoNet [8], appears to have been terminated in 2017. For some research areas existing testbeds have network configurations that are too dense. An example is Source Location Privacy (SLP) [6], where a WSN is intended to be deployed over a large outside area to monitor valuable assets in which the network should not reveal the asset's location by the path messages take. A sparse network is required because some techniques (such as [5]) require sufficient spatial redundancy to provide effective privacy. One option to reduce the density in a network is for nodes to be turned off. Alternatively, by reducing the transmit power of devices, different network topologies can be obtained while still using all nodes in the network. The change in transmission power is often undertaken without consideration of the environment it exists in, and the effects this has on the performance of protocol running on the network. In other words, no benchmark exists to assess the performance of transmission power on the WSN.

Recent work, such as the IoT Benchmarks Initiative aims to quantify the performance of IoT, CPS and WSN systems by developing a standard way to test and compare wireless networking protocols [4]. Therefore, to work towards understanding the effects of approximating conditions for testing SLP protocols, in this paper we present a benchmark of TelosB nodes on FlockLab performed before we deployed SLP experiments on FlockLab. These results were used to understand the characteristics we could expect from the testbed. Based upon our experience, we also identify additional aspects of the testbed that needs to be monitored: (i) clock deviation of the controller nodes — to ensure the correctness of results, and (ii) availability of the testbed nodes — for repeatable experiments.

## 2 METHODOLOGY

A benchmark for a computing system is a set of programs for which well defined metrics exist that capture their performances and against which the performance of other programs are compared. As such, repeatability of performance measurement is important. For a WSN, developing a set of protocols for benchmarking a given WSN testbed is important. However, unlike benchmarking of traditional computing systems such as cluster systems, WSNs are susceptible to environmental perturbations. Further, different programs may require different network properties, such as network density.

Towards the development of such a benchmark, we focus on three specific environmental factors that will affect the benchmarking process. We focused on gathering performance results about



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 978-1-4503-6693-9/19/04...\$15.00  
<https://doi.org/10.1145/3312480.3313171>

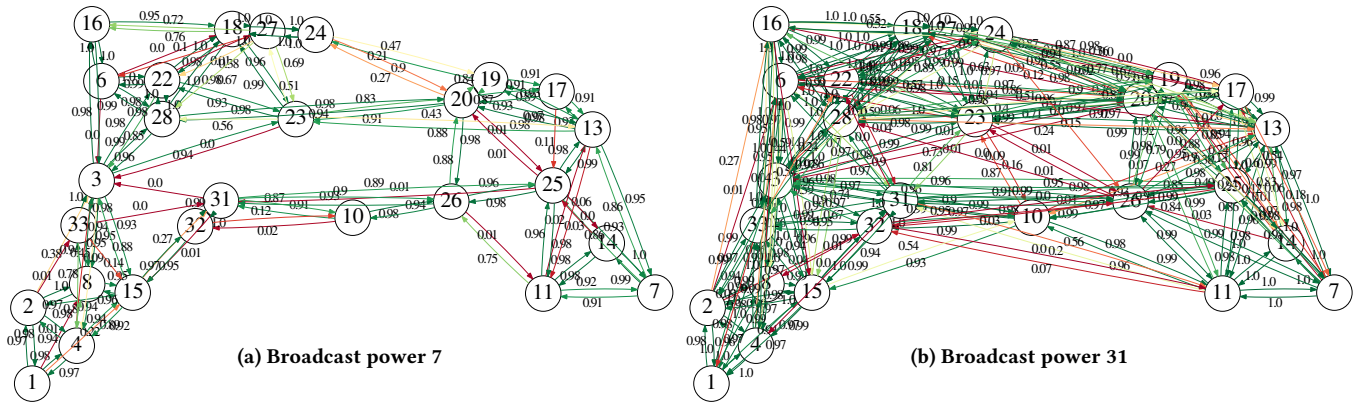


Figure 1: The probability of a message being delivered along the links between nodes on FlockLab.

three key areas: (i) the noise floor, (ii) the link performance, and (iii) the power consumption. The first two were important as they allowed noise and connectivity models to be created for simulators (such as TOSSIM [9]). The second also indicates the impact power level of transmissions has on topological control. The third is important to understand the minimum and maximum power consumption on the testbed, as this allows the power consumption of new protocols to be situated in the context of the testbed.

### 2.1 Measuring Noise Floor

To obtain noise floor measurements all TelosB nodes in the testbed were reserved. Each node continuously queries the background noise. Every 128 reads, the minimum, maximum and average background noise is sent over serial output to be logged. The implementation which measures the RSSI floor is based off the TinyOS RssiToSerial App<sup>1</sup>.

### 2.2 Transmit and Receive Performance

To obtain transmit and receive performance, an experiment is performed for every node in the network, where only that node broadcasts a message every 0.5 s and all other nodes listen for it. This is repeated for each node in on the testbed, at three transmit powers {7, 19, 31}. The nodes do not duty cycle and listen continuously. When a message is received the CC2420 chip is queried to obtain the received signal strength indicator (RSSI) and the link quality indicator (LQI). The number of messages sent is recorded and used to calculate the packet reception rate (PRR). We focus on these three metrics, however, other link quality metrics may also be used [3]. For example, signal to noise ratio (SNR) can be calculated from RSSI and noise floor, but raw RSSI values are useful as they can be provided to simulators. All testing was performed on channel 26 for reasons of space, but testing on other IEEE 802.15.4 channels (and using TSCH [13]) is important for understanding link performance.

A rate of 1 message every 2 seconds was chosen as it is not the maximum rate of the nodes, meaning that packet losses will be due to the radio environment and not hardware limitations. Applications with high data rates will need to investigate testbed performance when messages are sent without delay.

<sup>1</sup><https://github.com/tinyos/tinyos-main/blob/master/apps/tests/cc2420/RssiToSerial>

Power Level	Output Power (dBm)	Current Consumption (mA)
31	0	17.4
27	-1	16.5
23	-3	15.2
19	-5	13.9
15	-7	12.5
11	-10	11.2
7	-15	9.9
3	-25	8.5

Table 1: CC2420 Power levels [12, sec 28]

### 2.3 Power Consumption

One of the key advantages of the FlockLab testbed is its ability to measure the power consumption of the TelosB hardware with no additional work on the part of the user other than requesting it in the job’s configuration. To understand the energy cost of the algorithms being tested, it is important to understand the best and worst case energy consumption when performing different activities. As broadcasting at a lower transmit power will consume less energy (see Table 1), it will be important to understand how lowering the transmit power to obtain a sparse network will affect the power consumption.

## 3 RESULTS

In this section we present results from profiling the FlockLab testbed. The raw results are available from<sup>2</sup> and scripts to process these results can be found at<sup>3</sup>.

### 3.1 Noise Floor

Dips in the noise floor were observed at channels 11, 15, 20, 25 and 26 as these channels are in-between the three 2.4 GHz IEEE 802.11 channels 1, 6 and 11. Some areas of the testbed (such as nodes 1, 2 and 4) receive less interference on these channels, indicating that they may be further away from a WiFi router compared to other nodes in the testbed. This pattern of noise would not be the same if the testbed was located outdoors. If the WSN was located in a busy area, this noise would likely be greater. On the other hand, if located

<sup>2</sup><https://doi.org/10.5281/zenodo.2528758>

<sup>3</sup><https://bitbucket.org/MBradbury/slp-algorithms-tinyos/src/default/scripts>

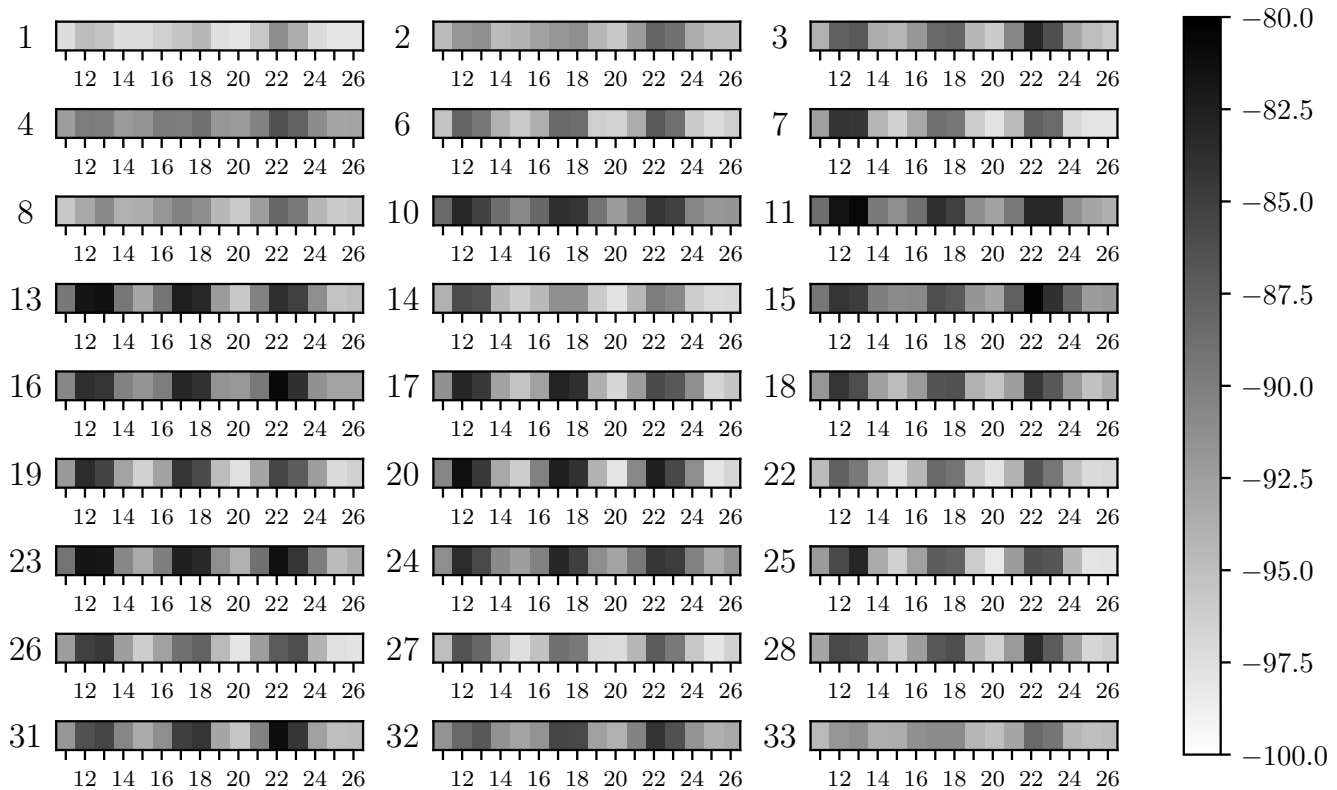


Figure 2: Noise floor (dBm) readings for FlockLab nodes on IEEE 802.15.4 channels 11–26.

in an area such as a forest, the noise would be lower. Techniques intended to be deployed outside will need to consider that the noise environment will be different to testbeds based indoors and also consider how this impacts the performance.

### 3.2 Transmit and Receive Performance

In Figure 1 and Figure 3 the results for PRR at different transmit powers on channel 26 is shown. The transmitting nodes are shown on the left and the receiving nodes on the bottom of the heatmap. No measurements are recorded between a node and itself. These figures demonstrate the effectiveness at lowering the transmit power to obtain a sparse network topology. To obtain a sufficiently sparse network with good link connectivity, the transmit power needed to be lowered to near the minimum (CC2420 transmit power level of 7). At this lower power level the majority of links have a high PRR, but there are also links with medium and low PRRs, which provides an interesting sparse topology to test on.

As well as PRR, both RSSI and LQI need to be considered together to understand the link performance. An LQI of near 110 is the maximum the CC2420 chip can produce and the minimum is near 50 [12, p. 48]. These are values that can be provided to TOSSIM or COOJA in order to simulate FlockLab before deployment.

### 3.3 Link Asymmetry

The degree of asymmetry of links is shown in Figure 4, with the colours indicating the degree of asymmetry where red and blue

indicate asymmetric links and white indicates symmetric links. It is important to test algorithms on networks with unidirectional links, as these can cause impossibility conditions [7]. FlockLab produces different levels of link asymmetry at different power levels, with links of different strengths at lower transmit powers, highlighting the challenge in developing benchmarks for WSNs and IoT.

### 3.4 Power Consumption

The average current draw (plus standard deviation error bars) is shown in Figure 5 for four scenarios:

- (1) When the nodes are performing no actions and the radio was not turned on (Nothing)
- (2) When the radio is on and it continuously polls for background noise (RSSI)
- (3) When a nodes transmits a packet every 0.5 s (the radio does not sleep) (Tx)
- (4) When nodes are listening for packets (Rx)

The current can be converted into power by multiplying by the voltage specified in the FlockLab configuration files (3.3 V). This is not done for these results, as the power profiling performed by FlockLab does not log the voltage, meaning we cannot be certain of the voltage at a specific time.

Figure 5 shows that current draw when performing the same task can vary by up to nearly 0.95 mA across different nodes. The cause of this could be due to a number of factors including: (i) poor calibration of current meters, (ii) environmental conditions,

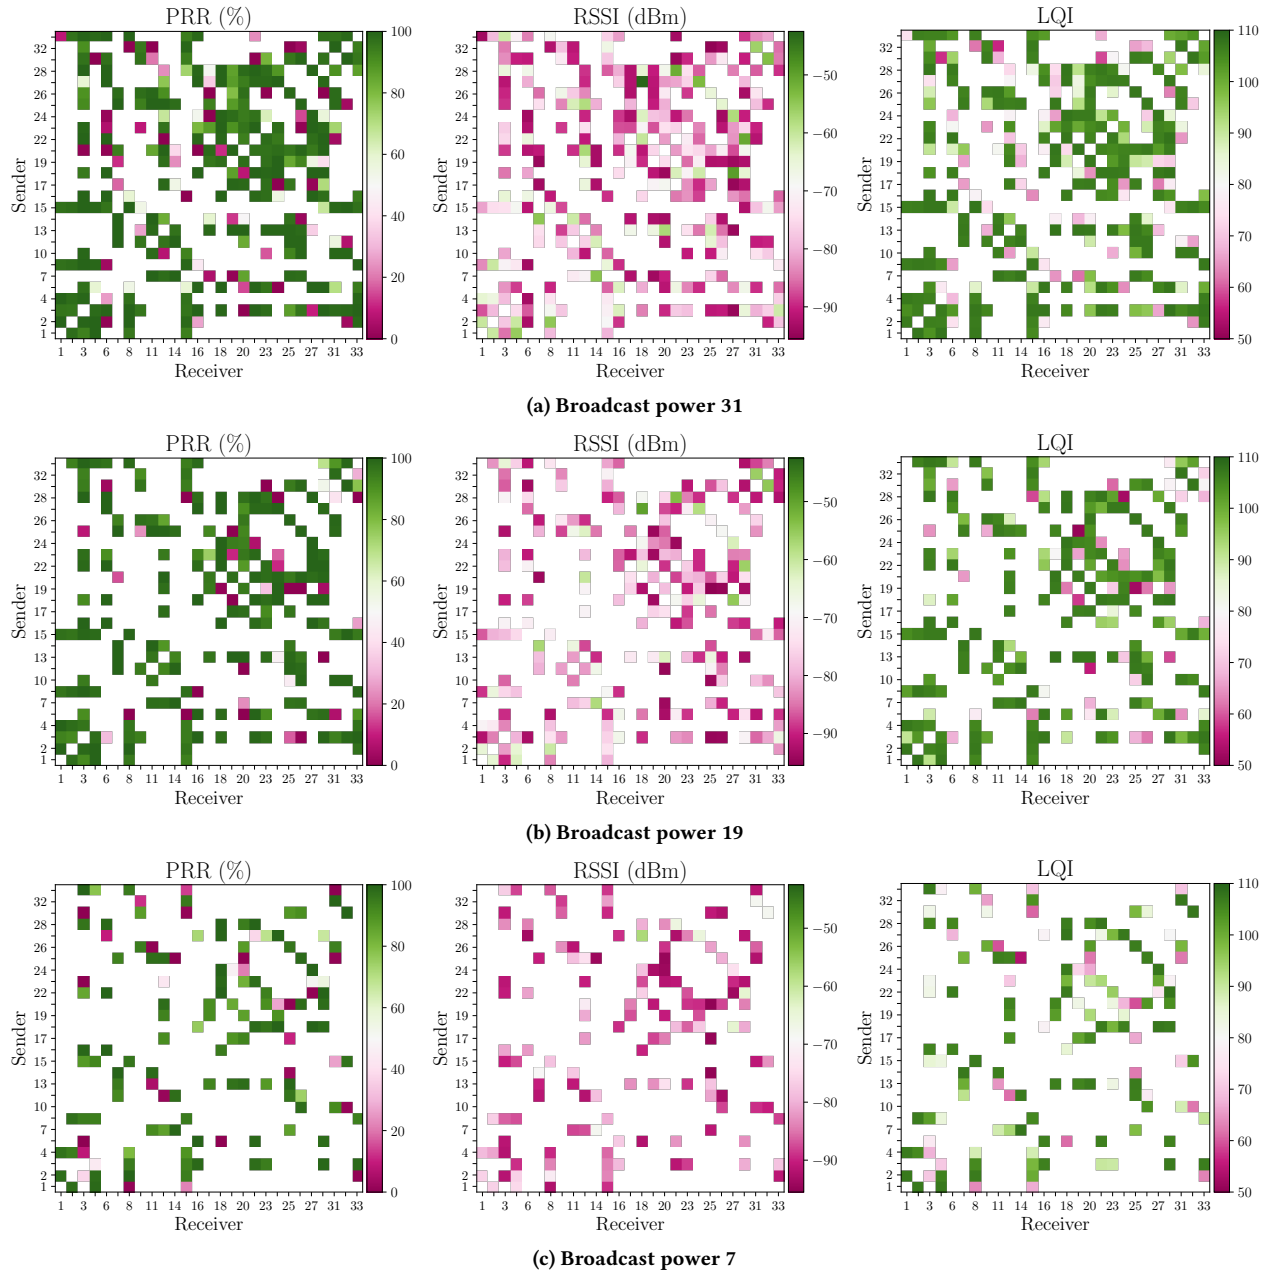


Figure 3: Link quality results for FlockLab on channel 26.

(iii) power supply issues, and other problems. This means that to perform a comparison of power consumption between nodes these variances need to be measured and accounted for in the analysis.

The second set of results, where the background noise is polled has the highest average power consumption. This is because the nodes are always active and the CPU has no opportunity to go into a lower power state. Once an RSSI value is obtained, this information is sent over the serial connection and then the process repeats without any delay. The importance of this is that logging via the serial port will have an effect on the current consumption.

Larger error bars are present when performing Tx/Rx as there are periods of high energy activity between sleeping. Nothing and RSSI consistently perform the same task, so have small error bars.

#### 4 EXPERIENCES USING FLOCKLAB

This section contains a discussion of the experiences gained from performing experiments on the FlockLab testbed. These experiences demonstrate the need for testbeds to be benchmarked and monitored in terms of operational aspects.

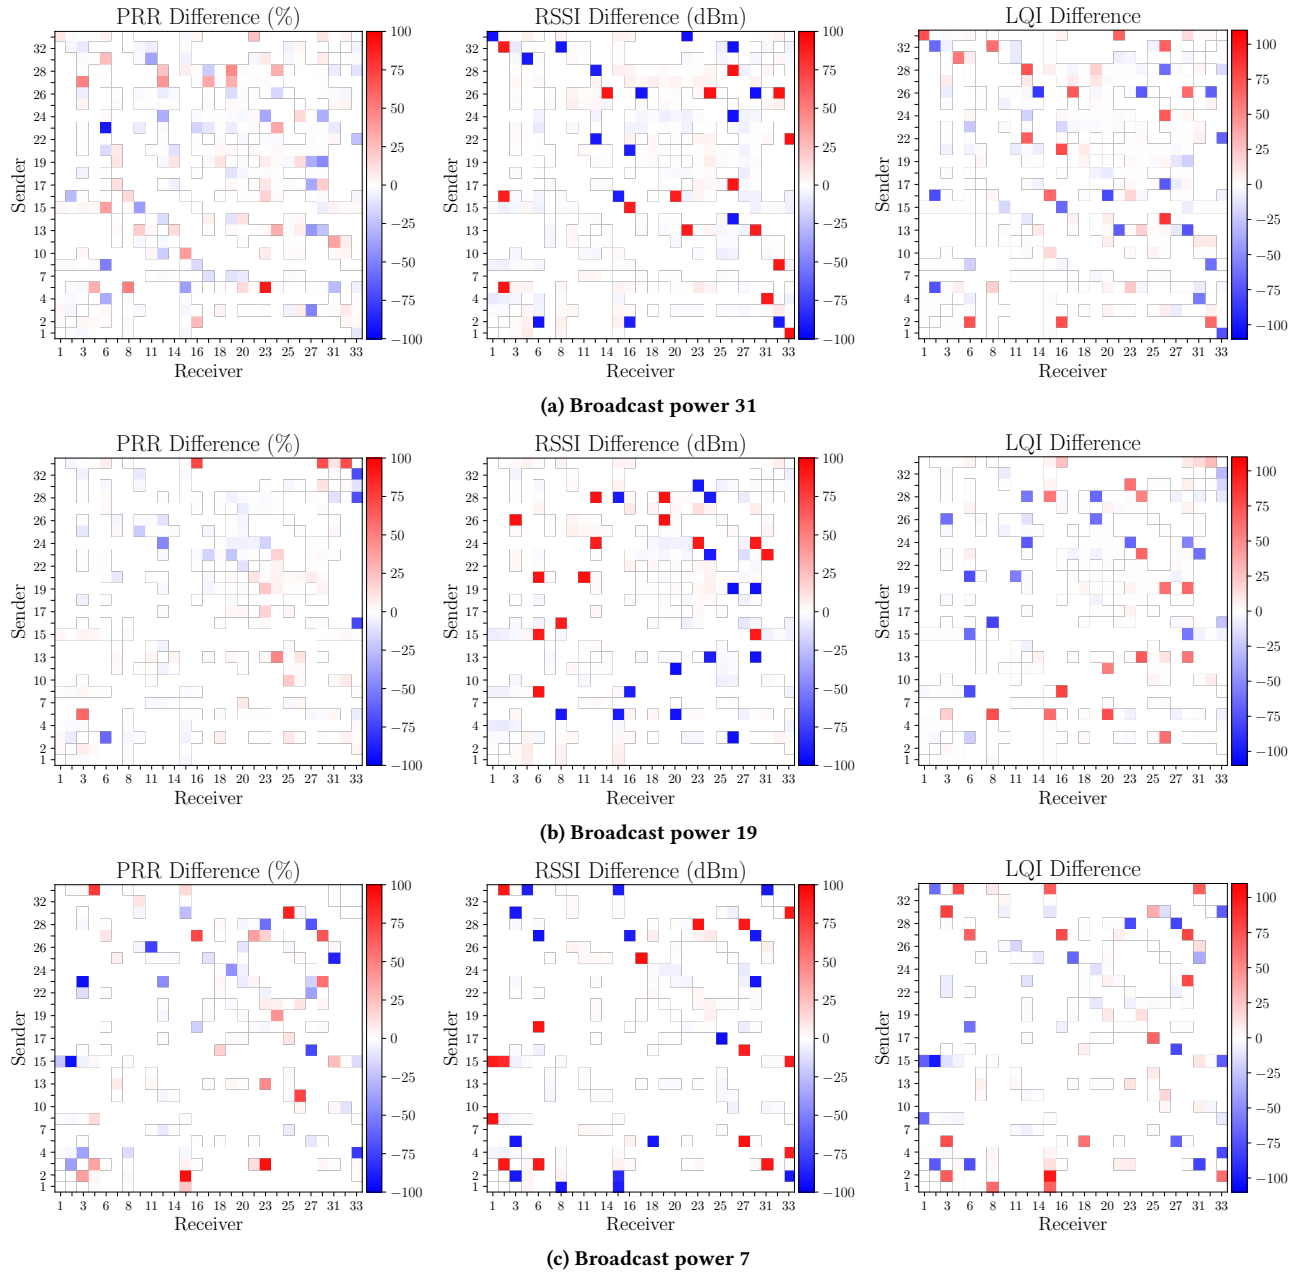


Figure 4: Link asymmetry results for FlockLab on channel 26.

#### 4.1 Time Synchronisation of the Global Testbed Clock

While deploying algorithms on FlockLab, an issue we encountered was that time was not correctly synchronised between nodes<sup>4</sup>. This issue was corrected by using a more accurate time server<sup>5</sup>. The result of this was that some experiments produced results stating impossible conditions (such as nodes receiving messages before

a neighbour had sent them). In order to compare the chronology of events that occur on different nodes, a consistent global clock needs to be available. This means that the timing of events on the infrastructure managing the testbed also needs to be benchmarked. Alternatively, techniques such as logical clocks (e.g., a vector clock) [11] could be used to record the order of events and thus avoid clock drift and synchronisation issues.

<sup>4</sup><https://groups.google.com/forum/#!topic/flocklab-users/INFfka-J5xQ>  
<sup>5</sup><https://gitlab.ethz.ch/tec/public/flocklab/commit/786ad707>

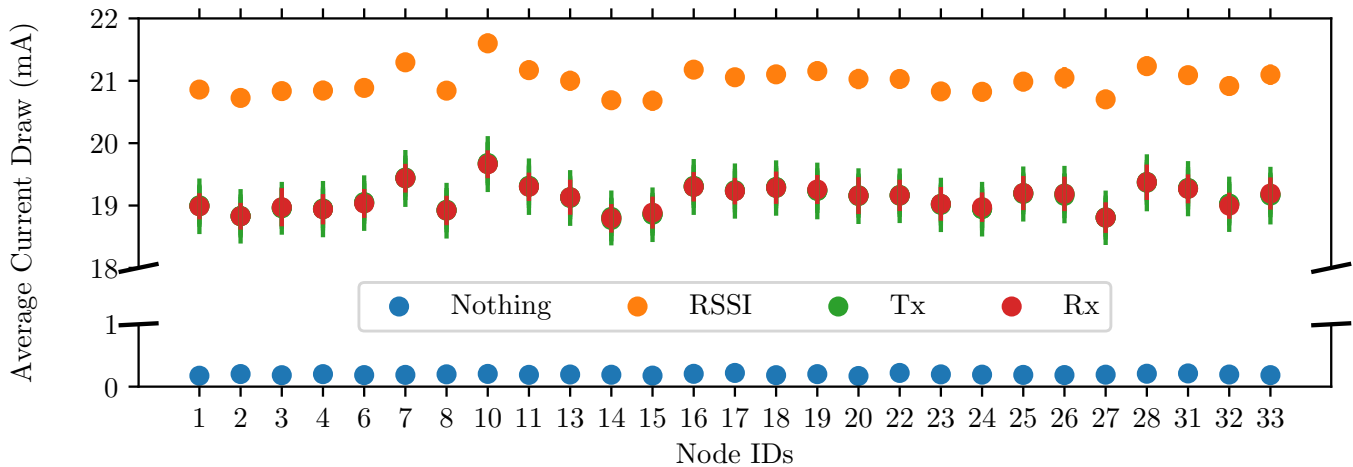


Figure 5: Average current draw (mA) in four different situations.

## 4.2 Node Availability

Whilst running other experiments nodes did not have a 100% uptime, which meant that certain nodes were excluded from tests to continue experiments. To ensure experiments cannot request unavailable nodes, the FlockLab submission script prevents a job from being submitted to the testbed if one of the nodes requested is unavailable. For experiments to be reproducible the same set of nodes will need to be requested each time. Consistent node availability was also an issue when using FIT IoT-LAB. This means that the testbed's reliability (each node's uptime) should be monitored. A testbed with low stability in terms of which nodes are available makes performing equivalent experiments difficult.

## 5 CONCLUSIONS

To conclude, decreasing the transmit power of TelosB nodes in the FlockLab testbed is an effective way to obtain a sparse network topology for testing. The results indicate that there is minimal impact on the average current draw, meaning that power measurements will still be useful. To be able to perform practical experiments of techniques that solve problems such as SLP, there is still a need for large-scale testbeds located outside, instead of an indoors environment. There is also a need to monitor operational aspects of the testbed – including time synchronisation and availability of nodes – to ensure that results are accurate and repeatable. It would also be useful if testbeds performed this kind of profiling while no jobs are running, to provide an up-to-date state of the testbed.

## ACKNOWLEDGMENTS

This research was supported by the Engineering and Physical Sciences Research Council (EPSRC) through [DTG Grant EP/M506679/1] and the PETRAS: UK Research Hub for Cyber Security of the Internet of Things project [EPSRC Grant EP/N02334X/1].

## REFERENCES

[1] 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>

[2] P. Appavoo, Ebram K. William, M.C. Chan, and M. Mohammad. 2018. Indriya2: A Heterogeneous Wireless Sensor Network (WSN) Testbed. In *TRIDENTCOM*. <https://indriya.comp.nus.edu.sg/Indriya2.pdf>

[3] N. Baccour, A. Koubâa, L. Mottola, M. A. Zúñiga, H. Youssef, C. A. Boano, and M. Alves. 2012. Radio Link Quality Estimation in Wireless Sensor Networks: A Survey. *ACM Trans. Sen. Netw.* 8, 4, Article 34 (Sept. 2012), 33 pages. <https://doi.org/10.1145/2240116.2240123>

[4] C. A. Boano, S. Duquennoy, A. Forster, O. Gnawali, R. Jacob, H. Kim, O. Landsiedel, R. Marfievici, L. Mottola, G. P. Picco, X. Vilajosana, T. Watteyne, and M. Zimmerling. 2018. IoT-Bench: Towards a Benchmark for Low-power Wireless Networking. In *Proceedings Of The First Workshop On Benchmarking Cyber-physical Networks And Systems (cpsbench 2018)*.

[5] M. Bradbury, A. Jhumka, and M. Leeke. 2018. Hybrid Online Protocols for Source Location Privacy in Wireless Sensor Networks. *J. Parallel and Distrib. Comput.* 115 (May 2018), 67–81. <https://doi.org/10.1016/j.jpdc.2018.01.006>

[6] M. Conti, J. Willemsen, and B. Crispo. 2013. Providing Source Location Privacy in Wireless Sensor Networks: A Survey. *IEEE Communications Surveys and Tutorials* 15, 3 (2013), 1238–1280. <https://doi.org/10.1109/SURV.2013.011413.00118>

[7] A. Jhumka and L. Mottola. 2016. Neighborhood View Consistency in Wireless Sensor Networks. *ACM Trans. Sen. Netw.* 12, 3, Article 19 (July 2016), 41 pages. <https://doi.org/10.1145/2901296>

[8] U. Kulau, S. Schildt, S. Rottmann, B. Gernert, and L. Wolf. [n. d.]. Demo: PotatoNet – Robust Outdoor Testbed for WSNs: Experiment Like on Your Desk. Outside.

[9] P. Levis, N. Lee, M. Welsh, and D. Culler. 2003. TOSSIM: accurate and scalable simulation of entire TinyOS applications. In *Proceedings of the 1st international conference on Embedded networked sensor systems (SenSys '03)*. ACM, 126–137. <https://doi.org/10.1145/958491.958506>

[10] R. Lim, F. Ferrari, M. Zimmerling, C. Walser, P. Sommer, and J. Beutel. 2013. FlockLab: A testbed for distributed, synchronized tracing and profiling of wireless embedded systems. In *2013 ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. 153–165. <https://doi.org/10.1145/2461381.2461402>

[11] M. Raynal and M. Singhal. 1996. Logical time: capturing causality in distributed systems. *Computer* 29, 2 (Feb 1996), 49–56. <https://doi.org/10.1109/2.485846>

[12] Texas Instruments. [n. d.]. CC2420 Datasheet. Online. [ti.com/lit/ds/symlink/cc2420.pdf](http://ti.com/lit/ds/symlink/cc2420.pdf) Accessed: 21/09/2017.

[13] T. Watteyne, M. Palattella, and L. Grieco. 2015. Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement. RFC 7554. <https://doi.org/10.17487/RFC7554>

[14] X. Zhong, Y. Liang, T. A. Slater, F. Plaza, and X. Liang. 2018. Network Dynamics Analysis and Benchmarking on an Outdoor Heterogeneous Wireless Sensor Network. In *Proceedings of the 15th ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, Ubiquitous Networks (PE-WASUN'18)*. ACM, New York, NY, USA, 92–99. <https://doi.org/10.1145/3243046.3243060>