1KT: A Low-Cost 1000-Node Low-Power Wireless IoT Testbed

Mateusz Banaszek, Wojciech Dubiel, Jacek Łysiak, Maciej Dębski, Maciej Kisiel, Dawid Łazarczyk, Ewa Głogowska, Przemysław Gumienny, Cezary Siłuszyk, Piotr Ciołkosz, Agnieszka Paszkowska, Inga Rüb, Maciej Matraszek, Szymon Acedański, Przemysław Horban, Konrad Iwanicki
Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Poland

ABSTRACT

Testbeds remain indispensable instruments for experimentally evaluating IoT-oriented low-power wireless networking solutions. With the evolution of the field, they are increasingly expected to match envisioned deployment conditions of such solutions, notably in terms of scale. However, large-scale testbeds are scarce, likely because they have been believed to be expensive. This paper argues that this belief need no longer be justified by presenting the architecture and basic properties of 1KT, our new smart-building IoT testbed for solutions utilizing IEEE 802.15.4 and Bluetooth Low Energy. It comprises 1000 experimental devices deployed directly in human spaces of 168 rooms on all 5 floors of a sizable building. At the same time, its cost is relatively low considering the scale.

CCS CONCEPTS

• **Networks** → **Sensor networks**; **Network experimentation**; • **Hardware** → **Sensor applications and deployments**; **Sensor devices and platforms**.

KEYWORDS

1KT, testbed, experiment, experimental testbed, large scale, low-power wireless network, Internet of Things (IoT), smart building, IEEE 802.15.4, Bluetooth Low Energy (BLE)

ACM Reference Format:
Mateusz Banaszek, Wojciech Dubiel, Jacek Łysiak, Maciej Dębski, Maciej Kisiel, Dawid Łazarczyk, Ewa Głogowska, Przemysław Gumienny, Cezary Siłuszyk, Piotr Ciołkosz, Agnieszka Paszkowska, Inga Rüb, Maciej Matraszek, Szymon Acedański, Przemysław Horban, Konrad Iwanicki. 2021. 1KT: A Low-Cost 1000-Node Low-Power Wireless IoT Testbed. In *Proceedings of the 24th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '21), November 22–26, 2021, Alicante, Spain.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3479239.3485708

1 INTRODUCTION

Low-power wireless networking, notably utilizing IEEE 802.15.4 and Bluetooth Low Energy (BLE), is one of the key technologies for the Internet of Things (IoT). However, communication enabled by these two narrowband 2.4 GHz standards is peculiar in several aspects. For instance, it proceeds over wireless links that are heavily influenced by the ambient conditions, like temperature [5], and various obstacles, notably humans [16]. Interference, from both coexisting links and other technologies, poses yet another challenge, especially for the ubiquitously used 2.4 GHz band. Although many of such problems have been studied extensively and compelling solutions have been proposed by the community to alleviate them, the field has recently started facing a fundamental methodological challenge: our abilities to evaluate the performance of the proposed solutions increasingly fail to meet their prospective industrial adopters' expectations, which is a major impediment also to further scientific progress [4].

To explain, given the inherent limitations of today's simulators in accurately yet efficiently modeling low-power wireless networking at the lowest levels [4, 13], novel communication schemes have to be extensively evaluated also on testbeds. One approach is to isolate such testbeds from their surrounding environment as much as possible and to emulate external radio interference, human presence, and other relevant phenomena on demand. Such highly controllable settings are indispensable especially for basic research at early, small-scale stages of protocol development. However, given the peculiarities of low-power wireless communication, experiments in more realistic conditions often reveal unpredicted protocol behaviors, like performance deteriorations [6, 9] or, even more surprising, improvements [14]. For these reasons, from an industrial perspective, especially when dependability is a core requirement [11], a protocol considered for real-world deployments should first be evaluated in settings similar to the envisioned ones. Such results can be obtained on testbeds that model as close as possible the target conditions by being themselves located in similar environments and having matching scales. Appropriate low-power wireless networking testbeds are scarce, though, likely because developing such testbeds has simply been believed as overly expensive.

In this paper, we argue that this belief need no longer be justified and support the claim by introducing *1KT*, our new testbed dedicated to IoT-oriented networking research. It offers for experimentation 1000 custom devices, featuring modern ARM Cortex-M3s, radios for both IEEE 802.15.4 and BLE, and a range of other functionalities, like continuous power consumption measurement. It covers all floors of a sizable five-story faculty building, with the devices deployed directly in human spaces (e.g., under windowsills, behind desks, on file cabinets) of 168 rooms of virtually all types, including offices, computer labs, lecture halls, and shared areas. In effect, the low-power wireless network they form is volumetric, nonuniform, and has a large diameter. Moreover, it is subject to multiple short- and long-term phenomena that can be related to specific events and processes occurring at the deployment site and stemming from the intended building functions, which constitutes yet another evidence for the need for experimentation on such platforms. At the same time, the cost of 1KT is relatively low considering its scale, which refutes the common presumption that a testbed, especially a large-scale one, requires an immense budget.

MSWiM '21, November 22–26, 2021, Alicante, Spain.

ACM ISBN 978-1-4503-9077-4/21/11.
https://doi.org/10.1145/3479239.3485708

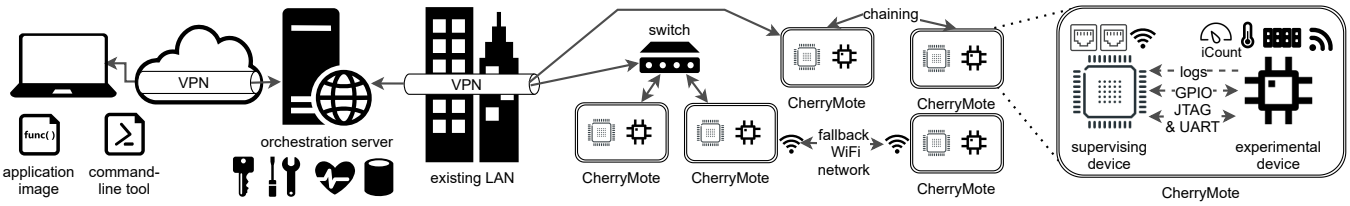

Figure 1: System architecture

The rest of this paper is organized as follows. We start by discussing the design and implementation of 1KT (Section 2). Then, we empirically assess its basic properties (Section 3). Finally, we survey related work (Section 4) and conclude (Section 5).

2 DESIGN AND IMPLEMENTATION

Our testbed adopts a popular architecture that combines: experimental devices, a control network, and an orchestration server (cf. Figure 1). Yet, to reconcile the goal of the large scale with the limited budget, its implementation requires trade-offs, such as limiting the offered features to essential ones only.

At a high level, 1KT follows three primary decisions. First, it is homogeneous to benefit from the effects of scale and to simplify deployment and maintenance. Second, it is built out of off-the-shelf components to be cost-effective. Finally, special care is committed to the control network as non-remote operation and maintenance are impossible when 1KT is already deployed.

The Experimental Devices (EDs). If any technology is worth investing in, it is certainly EDs, as they execute the experiments aiming at studying particular phenomena of low-power wireless communication. Texas Instruments CC2650 is chosen to this end, as it is a modern industry-oriented SoC (Cortex-M3, 28 KB RAM, 128 KB Flash), with a built-in radio that supports both IEEE 802.15.4 and BLE. The choice of CC2650 is led also by its features that support the process of experimentation itself, like an integrated additional processor for offloading the CPU in some tasks. For instance, we take advantage of it to provide highly efficient logging for experimental applications running on 1KT. However, developing a custom CC2650-powered module would exceed the project's budget. Therefore, a CC2650EM-7ID Reference Design, which is a ready-to-use evaluation module, is employed directly as ED.

The Control Network – Communication Links. The control network allows the server and EDs to reliably exchange data and commands by means other than the low-power wireless communication that is the subject of the experiments. However, deploying a dedicated wired physical network would abuse the hospitality of the testbed host and unjustifiably raise costs. In contrast, relying solely on wireless technologies, whose quality may be affected by factors beyond our control, could impair the dependability of the testbed. Consequently, as the main medium for the control network, 1KT employs the Ethernet infrastructure already existing in the building, but augmented with virtualization technologies, notably VPN. As a result, however, the design of the testbed has to minimize the consumed bandwidth and employ additional Ethernet switches to prevent downgrading the network where the building occupants' computers are connected. As a fail-over communication medium, 1KT maintains a private Wi-Fi network.

EDs cannot be connected directly to the control network, though, as they lack proper interfaces and do not support remote supervision, and hence the need for extending the control network with devices that manage the communication links and operate EDs.

The Control Network – Supervising Devices (SDs). Such devices are combined one-to-one with EDs, as this approach has several advantages. First, it allows for placing both devices in a single case, which betters the overall appearance and simplifies deployment. Second, it improves reliability, in particular because a failure of SD can affect only one ED. On the other hand, the drawback of the selected approach is a potentially larger cost compared to an alternative scheme of managing a few EDs by one SD.

A single-board computer Olimex RT5350F-OLinuXino is chosen as SD because it is the most cost-effective option that meets our requirements. First, it supports low-level interfaces provided by ED, thereby facilitating integration. Second, it has a built-in Ethernet chip and Wi-Fi radio, so it can be configured to constitute a host in the control network. Moreover, it supports two Ethernet ports, allowing us to minimize the number of required Ethernet switches by *device chaining*: an SD relays network traffic to a connected SD.

Using economic hardware involves some trade-offs. For instance, RT5350F-OLinuXino features only 2.4 GHz Wi-Fi. Therefore, to avoid unwanted interference, experiments employing the low-power radios should be run on non-overlapping channels. Furthermore, the highly limited persistent storage of RT5350F-OLinuXino requires extending each SD with an external USB flash drive.

Regarding software, SD runs OpenWrt, a Linux-based OS, thereby empowering users with fine-grained control of their experiments, especially the ability to launch custom scripts and programs on SDs. At the same time, we provide a range of dedicated convenience services, such as automatic transmission of logs received from ED to the server. Moreover, the employment of well-known Linux-based solutions allows us to minimize the effort required to establish software that manages the control network.

When developing SD, a massive part of our efforts focused on ensuring high dependability of the remote access to SD (and thus to ED): 1KT was designed to troubleshoot potential problems entirely remotely. To minimize the risk of losing this functionality both standard and abnormal testbed operations had to be considered thoroughly. For instance, we ran multi-round tests to verify network-loss scenarios and upgrades of the entire software. We also aimed to envision users' mistakes and intentional hostile actions, and evaluated those scenarios as well. Lessons learned in

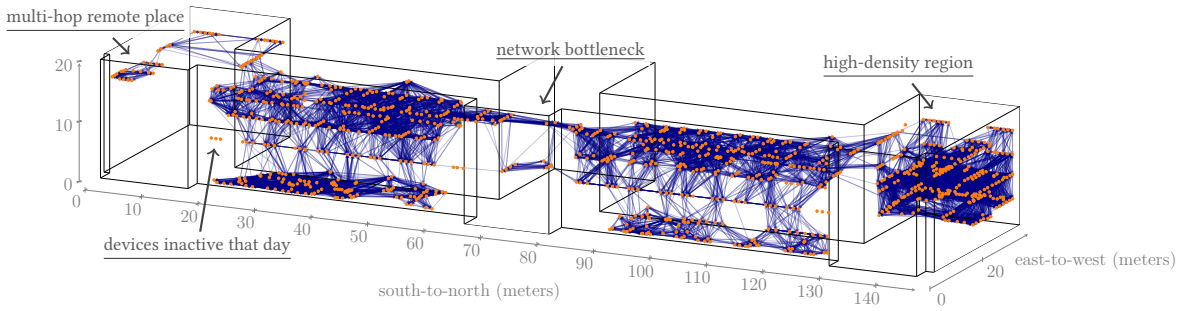

Figure 2: Visualization of links within the building established on a particular experiment day.

the course of hundreds of man-hours dedicated to ensuring reliability of 1KT taught us, however, that not every problem can be predicted [3]. Therefore we employed additional mechanisms enabling autonomous or remote recovery (e.g., a hardware-software watchdog or a last-resort remote service).

CherryMote – the Package. Each ED and its SD are coupled by being mounted on a PCB, which is the only component designed specifically for 1KT. Aside from connecting ED with SD, the PCB distributes power and provides additional connectors that allow for plugging external hardware or debugging equipment. To provide a way for each ED to interact with its surrounding environment, the PCB hosts a thermometer, a LED, and GPIO lines (connected to SD). The arguably limited set of sensors and actuators stems from the fact that 1KT is designed to support primarily low-power communication experimentation, and not to abuse the building's occupants' privacy. Moreover, each PCB hosts also a power meter (based on a design called iCount [7]), which can be queried in real time by both SD and ED itself.

All the aforementioned components are enclosed in a single case, constituting what we call CherryMote [3]. The choice of the particular professional yet off-the-shelf case is driven equally by purely technical requirements (protecting firmly all electronic components but being small enough to fit certain deployment locations like windowsills), and by visual taste (as the overall form of the device, we argue, is crucial for deployments in human spaces).

The Orchestration Server. The goal of the orchestration server is to coordinate all devices of 1KT so that they appear to their users as a single system. This is accomplished by two types of services. First, the server constitutes experimenters' gateway into the testbed, supporting launching and controlling experiments, as well as collecting and storing their results. Second, it drives the maintenance of the testbed, passing operators' commands to the devices, and providing real-time and historic information on their health. Such a centralized approach is fairly common among IoT testbeds, despite resulting in a single point of failure. Nevertheless, even in case of a fatal incident, the single server can be relatively easily replaced and restored from backups.

A user can access the testbed via a command-line tool, which can both be used to launch simple commands and be incorporated into scripts to execute more sophisticated experiments.

Maintenance and Costs. To date, virtually no significant maintenance tasks have been required. Despite some initial concerns, CherryMotes do not cause network problems nor do they interfere in any other intrusive way with the environment. Currently, although individual devices get disconnected from time to time in a partially uncontrolled manner, 1KT has been generally accepted by the building's occupants.

1KT was developed as an academic research project in which 18 people were involved (academic and administrative staff, graduate and PhD students). The project was granted 0.27M USD for 4 years.

Each complete and ready-to-use CherryMote costed around 85 USD. The deployment itself consumed about 1300 Ethernet cables, 400 power strips and 200 Ethernet switches, which amounted to about 11,000 USD for the whole 1KT. Physical arrangement of the testbed took our research team approximately 240 hours in total. It required neither specialized knowledge nor tools.

The only actual cost that has to be paid so far for operating 1KT is due to the consumed electricity, but it is relatively low: a single CherryMote draws around 2 Watts, a single Ethernet switch we used consumes around 1.5 W. Altogether, we estimate that the whole 1KT, including the server, requires around 2.5 kW.

3 EMPIRICAL PROPERTIES

Having the testbed fully operational, we ran a series of preliminary experiments to study its low-power wireless characteristic. In the following discussion, we present just their tiny subset. More specifically, we analyze basic connectivity properties of 1KT, and whether they are subject to changes in the environment, the knowledge of which is crucial given the purpose of the testbed.

In the experiments, each ED broadcast one IEEE 802.15.4 packet at a random moment in every fixed interval, nearly 300 packets per day. The communication proceeded on a commonly adopted channel 26, minimizing interference from other technologies. EDs were configured to communicate at presumably the most interesting transmission power value: the empirically identified minimal value for which all nodes constituted a single connected network (-3 dBm).

Each transmission and reception event was recorded. Then the data was processed to infer directed links between nodes (EDs), and compute their packet reception rate (PRR) on each day of every experiment. We focus on connectivity graphs induced by two classes of links depending on their PRR: links with PRR > 0% (i.e., those over which at least one packet was received), referred to as

all links, and links with PRR ≥ 90% (i.e., those of a high quality), dubbed *good links*. The initial analysis was based on data covering a full regular week. Throughout that time 1KT established over 60,000 *all links* and 30,000 *good links*.

Basic Properties. The data on the physical location of each device allows us to visualize the links (cf. Figure 2) and calculate their approximate lengths, if they were line-of-sight: an average *all link* measures 8 meters while a maximal—50 m, an average *good link* is 6 meters and a maximal—30 m.

The density of the low-power wireless network of 1KT can be estimated by the number of links coming into each node (in-degree) in the connectivity graph (cf. Figure 3a). The in-degree values of the *all links* graph are distributed over a relatively wide 5–175 range with the median of 60. In contrast, the distribution of in-degree over *good links* is concentrated in the 10–50 range.

A sequence of the minimal number of links (hops) a packet has to travel from its sender to its recipient, if multi-hop communication (e.g., routing) is employed, is called a minimal path (cf. Figure 3b). Over the course of the week, the longest minimal path (i.e., the diameter of 1KT) ranges from 10 to 12 hops over *all links*, and from 18 to 23 hops over *good links* only.

Environment Changes. The values presented hitherto are aggregated over one week of the experiment. Analyzing each day separately reveals noticeable changes in the wireless environment that occur between a workday and a weekend. For example, the number of *all links*, relatively constant over the workdays (around 60,000 links), drops below 50,000 during the weekend, whereas the number of *good links* raises over the weekend from 20,000 to around 30,000.

An analysis of link PRR in both directions unveils more details about the change with respect to quality and symmetry of the observed links. In general, most of the links are indeed symmetric, including a vast number of links with a high PRR (over 90%) in both directions and a significant number with a low PRR (near 0%).

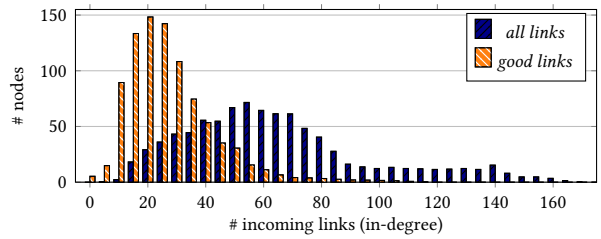

(a) Histogram of incoming links (in-degree) – neighborhood size.

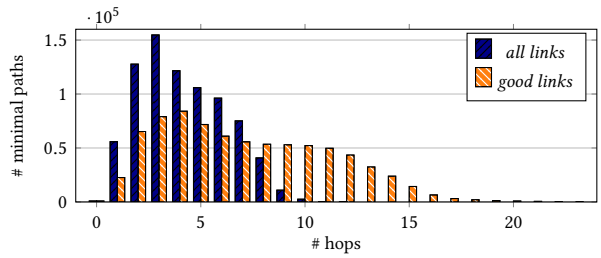

(b) Histogram of minimal path lengths between each ED pair.

Figure 3: Metrics averaged over all days of the experiment.

However, during a workday there are significantly more mid-PRR links with varying level of symmetry than over a weekend.

A possible reason of the change of the wireless environment between a workday and a weekend is that the physical environment is more stable during the latter: far fewer people and Wi-Fi-enabled devices in the building cause less interference but also fewer short-lived alternative paths of radio signals. Moreover, similar changes were observed both in a long term (links observed on workdays during a COVID-19 lockdown resembled the ones observed during the regular weekend) and a short term (local disturbances of the wireless network correlated with classes schedule).

Result Discussion. The results suggest that 1KT establishes an entire spectrum of links, in terms of PRR, a wide range of regions, in terms of density, and numerous multi-hop paths, making it particularly suitable for evaluation of indoor networking protocols. The stable and repeatable results of long-term experiments (not detailed in this paper) indicate that 1KT is indeed capable of providing scientific-grade data.

The close analysis of numerous metrics reveals that a low-power wireless network, such as the one of 1KT, undergoes multiple micro and macro-scale changes. In case of the experiment, the changes are highly correlated with the physical changes occurring at the deployment site that closely follow both regular processes typical of the faculty building and exceptional events. This observation supports our claim that industry-ready IoT solutions should be evaluated in conditions resembling their expected operational environments to prove that they remain performant under such local phenomena. As each target environment may display slightly different phenomena, the requirement for real-world evaluations supports the need for IoT testbeds being affordable for a variety of research groups.

4 RELATED WORK

Testbeds have long been recognized as indispensable instruments for studying low-power wireless communication. Many of them have been built to date (cf. surveys by Tonneau et al. [17] or Kim et al. [12]); many have also been decommissioned. As the field is maturing, new testbeds are needed to address novel challenges that emerge in the context of IoT. Prominent today's examples include testbeds associated with the IoTBench initiative [4], notably IoT-Lab [1], FlockLab 2 [18], Indriya2 [2], and platforms that facilitate protocol evaluation, like D-Cube [15] as well.

IoT-LAB is an open testbed originally featuring over 2,720 devices of various types and spreading across 6 sites. The largest one, Grenoble, hosts over 600 experimental devices (as of May 2021 [10]), which is 40% less than 1KT. The devices are located in corridors of a single floor, in raised floors and dropped ceilings, following a regular scheme, unlike 1KT. A supervising device of IoT-LAB provides multiple advanced capabilities (e.g., sniffing and injecting wireless traffic). Based on the published schematics, we estimate the price of its components as approximately 95 USD, whereas the corresponding components of CherryMote amount to around 30 USD.[1] This might make the cost of replicating IoT-LAB Grenoble at the scale of 1KT too high for low-budget research groups.

[1]The comparison was made for a batch of 1000 devices using the Octopart BOM Tool (https://octopart.com/bom-tool) as of May 2021. The original BOMs were modified to substitute components that were no longer available.

FlockLab 2 is built of 30 supervising devices (as of May 2021 [8]) deployed indoors on a single floor of an office building and outdoors (at a few remote rooftop locations). Each supervising device can host up to 4 experimental devices that can be of different types. This approach might be particularly beneficial for heterogeneous testbeds, as it reduces the total cost per experimental device and simplifies management, but we do not find it advantageous in the case of homogeneous testbeds like 1KT. Based on the published bill of materials, we estimate the price of the supervising device's components to around 120 USD, which presumably results from its exceptional features (e.g., high fidelity tracing and GNSS-based time synchronization).

Indriya2 is deployed across 3 floors of a university building, with the experimental devices placed mostly in hard-to-reach places like ceilings. It originally featured 102 experimental devices, including 28 based on CC2650, which is the same microcontroller as in EDs of CherryMotes. They are supervised by mini desktops, to which they are connected via USB cables and hubs: a dozen devices per mini desktop. While this approach reduces the total cost of required supervising devices, it limits the supervising capabilities to mostly only serial I/O. Moreover, in case of a smart-building deployment targeting tens of rooms, it would introduce additional burden during deployment and would require routing cables in-between rooms, which might not be acceptable by the building owner.

To sum up, each of the aforementioned testbeds is a unique experimental platform. They differ in scale, the environment they are deployed in, the choice of IoT technology that can be put under test, the level of control of the experimental process, insight into executed experiments, etc. We do not aim to surpass or replace any of them as our work has an orthogonal goal: exploring how to deliver a large-scale testbed that has the potential to meet the industrial expectations regarding empirical protocol evaluation at costs that are affordable to many individual research groups.

5 CONCLUSIONS

Testbeds are vital for research and development of low-power wireless networking, which is in turn one of the building blocks for IoT. This need is particularly reflected in the requirements of industrial technology adopters, who expect novel IoT solutions to have their performance confirmed in conditions and at scales resembling intended target deployments. However, testbeds meeting these requirements, for instance, aiming at large smart buildings, are scarce, most likely because they have been presumed to be intricate and expensive. Our work suggests that this presumption need no longer be justified. To this end, it demonstrates that a 1000-node indoor testbed, which establishes a large multi-hop network featuring a full spectrum of links and regions with different characteristics, as well as exhibiting various temporal phenomena due to its environment, can be accessible to multiple groups worldwide. We thus hope that our work will facilitate evaluating novel low-power wireless networking solutions at scale and inspire development of new testbeds in a range of interesting settings. Accordingly, we are open to collaboration on such projects. More information about the testbed can be found on the project website:

https://www.mimuw.edu.pl/~iwanicki/projects/heni/1kt.html

ACKNOWLEDGMENTS

The authors would like to thank the numerous people at the University of Warsaw, notably the authorities and administrative staff, whose support has made 1KT possible. The presented experimental research was sponsored mainly by the National Science Center (NCN) in Poland under grant no. 2019/33/B/ST6/00448. The development of the testbed, in turn, was sponsored by the National Center for Research and Development (NCBR) in Poland under grant no. LIDER/434/L-6/14/NCBR/2015.

REFERENCES

[1] Cedric Adjih, Emmanuel Baccelli, Eric Fleury, Gaetan Harter, Nathalie Mitton, Thomas Noel, Roger Pissard-Gibollet, Frederic Saint-Marcel, Guillaume Schreiner, JulienVandaele, and Thomas Watteyne. 2015. FIT IoT-LAB: A Large Scale Open Experimental IoT Testbed. In *Proc. IEEE WF-IoT '15*. IEEE, 459–464.

[2] Paramasiven Appavoo, Ebram Kamal William, Mun Choon Chan, and Mobashir Mohammad. 2018. Indriya2: A Heterogeneous Wireless Sensor Network (WSN) Testbed. In *Proc. TridentCom '18*. Springer, Cham, 3–19.

[3] Mateusz Banaszek, Inga Rüb, Maciej Dębski, Agnieszka Paszkowska, Maciej Kisiel, Dawid Łazarczyk, Ewa Głogowska, Przemysław Gumienny, Cezary Siłuszyk, Piotr Ciołkosz, Jacek Łysiak, Wojciech Dubiel, Szymon Acedański, Przemysław Horban, and Konrad Iwanicki. 2020. Lessons from Communication Problems that Nearly Jeopardized Development of Hardware-Software Support for a 1000-Device IoT Testbed. In *Proc. EWSN '20*. Junction Publishing, 259–264.

[4] Carlo A. Boano, Simon Duquennoy, Anna Förster, Omprakash Gnawali, Romain Jacob, Hyung-Sin Kim, Olaf Landsiedel, Ramona Marfievici, Luca Mottola, Gian Pietro Picco, Xavier Vilajosana, Thomas Watteyne, and Marco Zimmerling. 2018. IoTBench: Towards a Benchmark for Low-power Wireless Networking. In *Proc. CPSBench '18*. IEEE, 36–41.

[5] Carlo Alberto Boano, Hjalmar Wennerström, Marco Zuniga, James Brown, Chamath Keppitiyagama, Felix Oppermann, Utz Roedig, Lars-Åke Norden, Thiemo Voigt, and Kay Römer. 2013. Hot Packets: A Systematic Evaluation of the Effect of Temperature on Low Power Wireless Transceivers. In *Proc. ExtremeCom '13*. ACM Press, 7—12.

[6] Thomas Clausen, Ulrich Herberg, and Matthias Philipp. 2011. A critical evaluation of the IPv6 Routing Protocol for Low Power and Lossy Networks (RPL). In *Proc. WiMob '11*. IEEE, 365–372.

[7] Prabal Dutta, Mark Feldmeier, Joseph Paradiso, and David Culler. 2008. Energy Metering for Free: Augmenting Switching Regulators for Real-Time Monitoring. In *Proc. IPSN '08*. IEEE, 283–294.

[8] FlockLab 2. 2021. *Wiki*. https://gitlab.ethz.ch/tec/public/flocklab/wiki

[9] Rodrigo Teles Hermeto, Antoine Gallais, and Fabrice Theoleyre. 2018. On the (over)-Reactions and the Stability of a 6TiSCH Network in an Indoor Environment. In *Proc. MSWiM '18*. ACM, 83–90.

[10] FIT IoT-LAB. 2021. *Project's website*. https://www.iot-lab.info/

[11] Konrad Iwanicki. 2018. A Distributed Systems Perspective on Industrial IoT. In *Proc. ICDCS '18*. IEEE, 1164–1170.

[12] Hiecheol Kim, Won-Kee Hong, Joonhyuk Yoo, and Seong-eun Yoo. 2015. Experimental Research Testbeds for Large-Scale WSNs: A Survey from the Architectural Perspective. *Int. J. Distrib. Sens. Netw.* 11, 3 (2015), 630210.

[13] Ivan Minakov, Roberto Passerone, Alessandra Rizzardi, and Sabrina Sicari. 2016. A Comparative Study of Recent Wireless Sensor Network Simulators. *ACM Trans. Sen. Netw.* 12, 3 (2016), 20:1–20:39.

[14] Inga Rüb, Maciej Matraszek, Piotr Konorski, Małgorzata Perycz, Aleksander Waśniowski, Dominik Batorski, and Konrad Iwanicki. 2019. 30 Sensors to Mars: Toward Distributed Support Systems for Astronauts in Space Habitats. In *Proc. ICDCS '19*. IEEE, 1704–1714.

[15] Markus Schuß, Carlo Alberto Boano, Manuel Weber, and Kay Römer. 2017. A Competition to Push the Dependability of Low-Power Wireless Protocols to the Edge. In *Proc. EWSN '17*. Junction Publishing, 54–65.

[16] Rahul C. Shah, Lama Nachman, and Chieh-yih Wan. 2008. On the Performance of Bluetooth and IEEE 802.15.4 Radios in a Body Area Network. In *Proc. BodyNets '08*. ICST, 1–9.

[17] Anne-Sophie Tonneau, Nathalie Mitton, and Julien Vandaele. 2015. How to choose an experimentation platform for wireless sensor networks? A survey on static and mobile wireless sensor network experimentation facilities. *Ad Hoc Networks* 30 (2015), 115–127.

[18] Roman Trüb, Reto Da Forno, Lukas Sigrist, Lorin Mühlebach, Andreas Biri, Jan Beutel, and Lothar Thiele. 2020. FlockLab 2: Multi-Modal Testing and Validation for Wireless IoT. In *Proc. CPS-IoTBench '20*. ETH Zurich, TIK.