Towards Cognitive Device Management: A Testbed to Explore Autonomy for Constrained IoT Devices

Anders Eivind Braten, Nattachart Tamkittikhun, Frank Alexander Kraemer and Doreid Ammar
Department of Information Security and Communication Technology
NTNU, Norwegian University of Science and Technology, Trondheim, Norway
{anders.e.braten, nattachart.tamkttikhun, kraemer, doreid.ammar}@ntnu.no

ABSTRACT
Providing constrained IoT devices with more intelligence is important to make them work optimally with regard to energy consumption and quality of data. To overcome the constraints of the sensors, we place cognition, i.e., learning and planning, in the cloud. In this demonstration paper, we present a testbed for exploring autonomy for constrained sensor nodes.

ACM Classification Keywords
K.6.4 System Management: I.2.11Distributed Artificial Intelligence

Author Keywords
Internet of Things; Solar Energy Harvesting; Autonomous Device Management; Machine Learning

MOTIVATION: AUTONOMOUS CONSTRAINED NODES
Miorandi et al. [2] emphazise the importance of including self-management and autonomic capabilities in devices connected to the Internet of Things (IoT). Often, the purpose of such devices is to collect as much data as possible, with the best possible quality. Since data collection requires energy, which is a scarce resource in most IoT systems, many devices harvest energy from their environment. However, they run the risk of depleting their battery if they lack knowledge of their energy budget, i.e., energy intake vs. energy consumption. This is a challenge since IoT devices often are too constrained in regard to both computation and communication to learn and plan ahead. Constrained devices usually also lack access to contextual data, like weather forecasts. Such data can contribute to better estimations. Outsourcing the estimation process to the cloud is a countermeasure to these constraints. We have built a lab to explore how to make constrained sensor nodes able to operate more autonomously. The lab has two goals: 1) to explore autonomous resource management for IoT devices using machine learning; and 2) to make the learning process itself autonomous.

IoT'17, October 22–25, 2017, Linz, Austria

ISBN 978-1-4503-5318-2.

DOI: 10.1145/3131542.3140282

Figure 1. Sensor nodes used for testing in real weather conditions.

SYSTEM OVERVIEW AND ARCHITECTURE
In our testbed, we use Libelium Waspmotes to collect data. They are connected to The Things Network (TTN) via LoRaWAN antennas. We have three distinguished sets of nodes, deployed in different ways. Some are placed in various locations in the city of Trondheim. These are used for long-term data collection. We also have equipment to do high-precision measurements of energy consumption in the lab. Our most prominent sensor nodes are the 8 solar powered sensor nodes we have deployed on the roof, as shown in Figure 1. These are used for experimentation and testing under real weather conditions. The sensor nodes provide sensor data such as CO_2 level, sound level and temperature, as well as energy-related meta data, e.g., battery level and solar charging current. The data is uploaded to TTN via an antenna located approximately 400 meters away. An overview of the architecture of the testbed is shown in Figure 2. We also collect the latest local weather forecast from the Norwegian Meteorological Institute. The collected data is stored as .csv files and then pre-processed in the cloud using Python and Pandas. Our ultimate goal is to build a system (using Scikit-Learn) that continuously learns, predicts and plans the energy budget for each device. To meet such requirements, we train a machine learning model every day. The model uses the most recent data (i.e., weather forecast and the position of the sun) as input to predict the energy intake per sensor node for the next 24 hours. Later, we use the predictions to identify the optimal sensing mode for each sensor device, i.e., instruct the sensor to send as much data as it can while keeping the battery level as stable as possible.

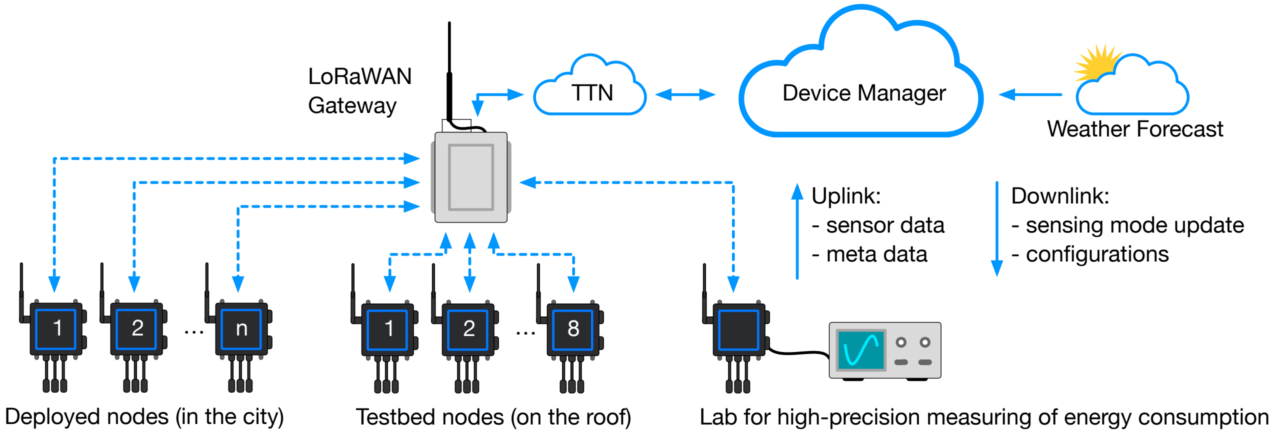

Figure 2. Deployment of the autonomous sensor testbed lab.

Through the TTN downlink, we instruct the sensor node to apply the identified sensing mode. The downlink is also used to update general configuration parameters, when necessary.

ENERGY AND SENSING CYCLE

The sensor node can adapt its own behavior by adjusting activities in different phases in its sensing cycle. We define the phases as sensing, data processing, transmitting data, and going to sleep for a period. The energy consumption of each phase can be adjusted by varying the execution time. This way, we can instruct the nodes to use different sensing modes to adjust their energy consumption behaviors. For instance, one sensing mode performs frequent sensing and calculates a mean to reduce the risk of errors, while another mode acquires the data sparingly to conserve energy. However, this comes with the tradeoff of lower data quality.

FINDINGS

By means of our testbed, we have made these discoveries:

- The energy consumption of applications running on a constrained device can be estimated with high precision using a generalized model of different sensing modes [3].
- Constrained devices can outsource research-intensive machine learning to a device manager located in the cloud. This is possible even when the communication channels themselves are constrained [1].
- To predict energy intake properly, it is important that the electronics of the sensor nodes reveal more data regarding the system operations, such as the current that is produced by a solar panel and the state of the charger [1].

To assist a system that also does the learning autonomously, we additionally need the following:

- Device management will have to manage, among other things, the process of knowledge gathering explicitly. This means that it also needs to take knowledge and learning into account, hence evolving into *cognitive device management*.
- To collect better training data, sensor nodes have to accept commands that override their normal operation. For instance, to gather more knowledge, they need to be able to switch on or off the charging, measuring short circuit currents, and switch to a different sensing mode on command. This way, the central device manager can choose to prioritize learning over normal operation in a period after a device has been deployed.

DEMONSTRATION SETUP

In the demonstration, we will: a) present an overview of our lab; b) give a real-time demonstration of how we are using the Waspmotes for noise detection; c) show interactive charts that visualize the collected data and the predictions that are made using machine learning; d) play time-lapse videos that shows the sensor nodes under real weather conditions.

ACKNOWLEDGEMENTS

We thank David Palma and Pål Sæther for their help in setting up the lab, and Dirk Ahlers and Fredrik Anthonisen for their support with the Waspmotes through the Carbon Track and Trace project.

REFERENCES

1. Frank Alexander Kraemer, Doreid Ammar, Anders Eivind Braten, Nattachart Tamkittikhun, and David Palma. 2017. Solar Energy Prediction for Constrained IoT Nodes Based on Public Weather Forecasts. In *The 7th International Conference on the Internet of Things (IoT 2017)*. ACM, Linz, Austria.

2. Daniele Miorandi, Sabrina Sicari, Francesco De Pellegrini, and Imrich Chlamtac. 2012. Internet of things: Vision, applications and research challenges. *Ad Hoc Networks* 10, 7 (2012), 1497 – 1516.

3. Nattachart Tamkittikhun, Amen Hussain, and Frank Alexander Kraemer. 2017. Energy Consumption Estimation for Energy-Aware, Adaptive Sensing Applications. In *International Conference on Mobile Secure and Programmable Networking (MSPN 2017) (LNCS)*, Bouzefrane et al. (Eds.), Vol. 10566. Springer, Paris, France, 247–260.