Utilising correlated information to improve the sustainability of Internet of Things devices

Jernej Hribar, and Luiz DaSilva
CONNECT - Research Centre
Trinity College Dublin, Dublin 2, Ireland
Email: {jhribar, dasilval}@tcd.ie

***Abstract*—Billions of low-power devices collecting information will be deployed in Internet of Things (IoT) networks. By taking advantage of the correlation exhibited in information collected, it is possible to improve sensors' energy efficiency. In this paper, we propose an updating mechanism capable of learning from the content of information collected to reduce the frequency with which devices transmit their updates, thus improving their energy efficiency. We show the potential gain of using correlated information and evaluate the proposed updating mechanism using data obtained from real sensing devices to determine the increase in energy efficiency.**

***Index Terms*—Machine Learning, Internet of Things, Age of Information, energy efficiency, low-power devices**

I. Introduction

Waldo Tobler, a famous geographer, once said "*Everything is related to everything else, but near things are more related than distant things*". The meaning of Tobler's first law of Geography is simple. The closer together two points on the map are, the more likely it is that the terrain between them remains the same. Similar logic applies to numerous Internet of Things (IoT) sensing devices collecting information essential for various services, e.g., traffic congestion control in smart cities [1], environment monitoring in smart factories [2], etc., to make decisions. The closer these devices are, the more likely it is that the collected information is correlated in time. Therefore, when information from a certain location is outdated or unavailable, correlated information from other IoT sensing devices may be used to aid services' decisions. Additionally, relying on available correlated information, an individual device can lower its transmission rate, without impacting the services' performance, and save energy. The main challenge with such an approach is to design an updating mechanism capable of effectively scheduling sensors' transmission times.

The decision of when an IoT device should transmit a new update depends on various factors, e.g., device's available energy, updates from other correlated sensors, contention for the transmission channel, etc. Additionally, low-power devices are constrained in available computing power and in the way they communicate with the rest of the network. The latter is especially noticeable if a device is battery-powered or relies on energy harvesting. In such a case, a device usually uses one of various Low-Power Wide-Area Network (LPWAN) standards [3] designed to support the device's use of sleep mode. That poses an additional challenge for the decision mechanism, as the decision has to be made while the device is still in active mode. For example, when a device is using LoRaWAN radio, whose message sequence we illustrate in Figure 1 [4], the network has only a short time-window available to make the decision. In such a network, the decision-making process has to adapt to the ever-changing environment constantly. To that end, we employ Reinforcement Learning (RL) [5] to design an adaptable updating mechanism.

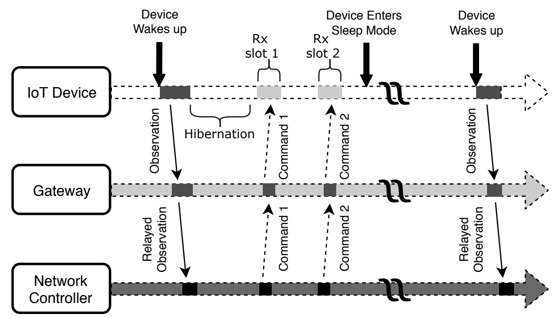

Fig. 1: Message sequence of an IoT device using LoRaWAN radio operating in class A.

In this paper, we propose an updating mechanism capable of decreasing the frequency with which IoT devices transmit updates without compromising the accuracy of the information collected. The mechanism learns from the content of information collected to improve the energy efficiency of low-power devices, thus making IoT deployments more sustainable, both economically and environmentally. In our work, the timeliness of information, i.e., the time elapsed since the last transmitted update, has a significant role in the decision-making process. We assess the timeliness of information using the concept of Age of Information (AoI) (Section III). Additionally, we demonstrate that relying on information from one correlated device can increase the accuracy of collected information on another device (Section IV). We describe how our proposed updating mechanism employs RL to leverage the exhibited correlation into energy savings. We evaluate our proposed updating mechanism using data collected in a real deployment (Section V). Finally, we discuss open issues and our future work (Section VI).

II. Related Work

The use of spatial and temporal correlation has inspired many energy saving schemes in the context of wireless sensor

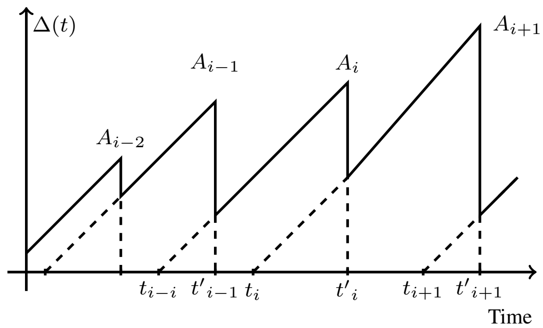

Fig. 2: Sample path for the process $\Delta(t)$.

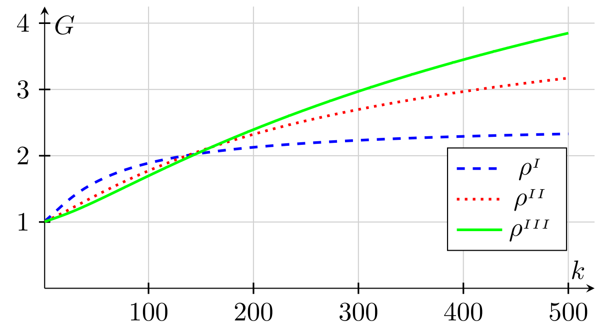

Fig. 3: The accuracy gain of first source as the second, correlated, source increases the frequency of updates. The selected covariance models, ρ^I, ρ^{II}, and ρ^{III}, are adopted from our analysis in [11]. The graph was obtained using covariance model parameters $a = 0.5$, $b = 0.025$, $c = 0.0125$, and $50s$ for the basic update time.

networks (WSNs). The authors of surveys [6] and [7] provide a thorough review of energy saving mechanisms for WSNs. Our work is most closely related to data reduction mechanisms, more specifically data prediction [8] and model-based active sampling [9] techniques. However, we are not relying on detection or reconstruction of the observed phenomenon to reduce the frequency of updates. Instead, we are proposing a mechanism employing RL to learn how to set sensors' transmission times according to environmental conditions. Additionally, we employ the idea of AoI, i.e., analyse the timeliness of updates, as presented in the next section.

III. Age of Information in the IoT

The concept of AoI was introduced in [10] to quantify the freshness of information as a metric of interest in vehicular networks. Information is always outdated due to delays in the network and, by measuring time elapsed since the information was generated, it is possible to make assumptions regarding its relevance. The fresher the information is, the more relevant it is for the decision making process. The *Age of Information* metric is used to determine the value of outdated information and to optimize the use of resources to deliver timely information.

An information source, e.g. a temperature sensor, generates packets with information regarding the observed physical phenomenon, e.g. temperature, and sends it to the sink, e.g., a thermostat. These packets of information are refereed to as status updates and contain a time stamp indicating the time of packet creation. The AoI is defined as the time that has elapsed since a status update was generated at the source:

$$\Delta_i(t) \coloneqq t - t_i, \tag{1}$$

where t_i represents the time the update was generated by the source. The behaviour of AoI at the sink is as presented in Figure 2. The value t'_i represents the time when the sink receives the status update. Whenever the sink receives the update, the AoI becomes $\Delta_i(t'_i) = t'_i - t_i$. The AoI of a status update always increases linearly with time until the next status update lowers the AoI value, at which point the AoI starts increasing linearly again.

The outdatedness of information impacts the accuracy of estimations the system can make regarding the observed phenomenon. Fresher information leads to more accurate estimates. By using updates from a correlated source, it is possible to prolong the times between two consecutive updates, i.e., increase the AoI of a status update, without lowering the accuracy of estimations. Next, we present the gain of using correlated information in a system of two correlated information sources.

IV. The Gain of Using Correlated Information

In a system of correlated information sources, the update rate and time difference between sources' updates impact the benefit of using correlated information. In [11] we analysed a system of two information sources. The system was interested in information collected by the first source, while the second source existed to assist the first source. When sources transmit with equal update times, the time between sources' updates plays a significant role in how beneficial the use of correlated information is. In such a case, the best updating strategy for the second source is to wait for the first source to transmit first before sending its update. We showed that there exist an optimal wait time for the second source for which the received correlated information will be the most beneficial for the system. When sources transmit with different update times, i.e., one source transmits more often, the system dynamics changes.

In particular, the case in which the second source transmits more often, which might happen if it has a lower cost of transmissions (e.g., has a more stable energy source), is most interesting. Building on the analysis we adopted in [11], we can show the increase of accuracy gain when using correlated information. We assume that the update rates between information sources are related by an integer factor k, meaning that the update rate of the second source is k-times higher than an update rate of the first source, i.e., the source of interest. We define the gain as the ratio between the estimation error when the system uses only updates from the source of interest and the estimation error when the system uses updates from both sources. Figure 3 shows that the more often the second source transmits, the higher the accuracy gain is.

To model the correlation between two sources, we used multiple covariance functions defined in [12] and [13]. Using

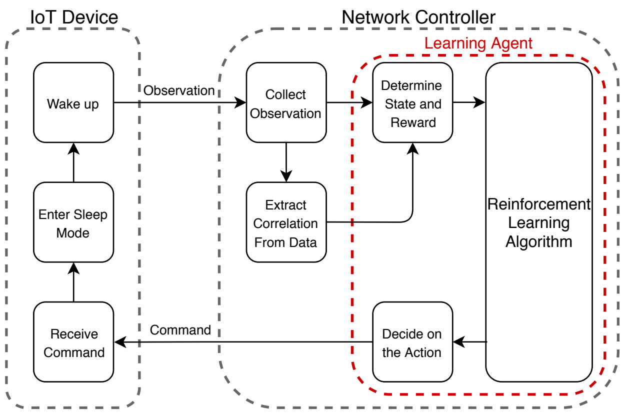

Fig. 4: Proposed updating mechanism.

multiple covariance models enables us to demonstrate the impact of spatial-temporal variation on the gain of using correlated information. A covariance function describes how observations from correlated sources jointly vary over time and space. For example, if sources measure temperature, the covariance reveals the probability that if temperature changed at the location of one source it has also changed at the location of the other. To make use of the demonstrated gains, we designed an updating mechanism capable of leveraging correlated information to prolong devices' lifetime, presented in the next section.

V. Proposed Updating mechanism

The updating mechanism needs to learn correlation from information collected and to be able to adjust devices' update time accordingly. The decision regarding the appropriate time between two consecutive updates for a specific sensor is based on multiple factors, e.g., the correlation between devices, available energy, etc. By setting devices' update time, the mechanism is effectively setting their AoI. Using RL enables the updating mechanism to find the optimal AoI, reflected in the update time, in a system with a changing environment. To apply RL to our problem we describe the system dynamics, i.e., devices' interactions with their surroundings, using a set of actions, states, and rewards.

Figure 4 shows a high-level model of the proposed updating mechanism. A transmission from an IoT device starts the learning cycle. In the first step, the network controller collects the transmitted observation. The observation is then used to estimate the correlation between devices, to determine device state, and to calculate the reward for the previously taken action. The state and reward depend on many parameters, such as energy level, the accuracy of observations, the frequency of transmissions, the action taken in the previous iteration, etc. The learning agent, depending on the state and reward, determines a new action. The learning agent's goal is to find a behaviour that will yield the highest long-term reward. The last step in the updating mechanism is to transmit the decision, i.e., the new update time, to the IoT device.

To evaluate the proposed updating mechanism, we use temperature measurements collected from 20 sensors deployed

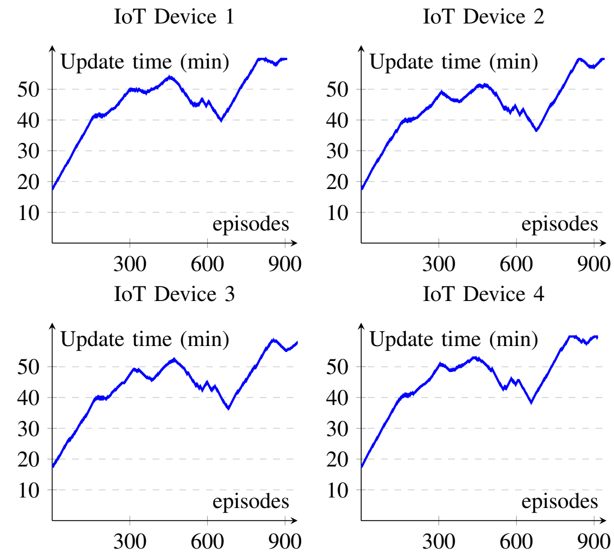

Fig. 5: Four IoT devices learning over 900 episodes.

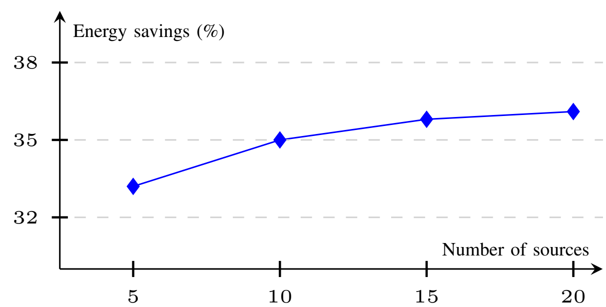

Fig. 6: Average energy saved per device.

in the SmartSantander testbed [14]. The data was gathered in April 2018. We implemented the mechanism using Deep Q-learning [15]. We define states and reward function the same as in our work in [16]. The system states depend on the device energy, update rate, and the accuracy of collected information. In this work, we have limited our available actions to three, all relative to the current device update time. The updating mechanism can increase or decrease the device's current time between updates for one time-step, or keep it as it is. We set the time-step to ten seconds. The system receives a positive reward if it has increased the time between updates and kept the accuracy within a specified range, or if it has decreased the sleep time when a higher rate of updates is required to preserve accuracy.

The primary goal of our mechanism is to increase the time between devices' updates if the system is collecting accurate information. As soon as the system notices that the information collected is no longer accurate, it will decrease the time between updates. Figure 5 shows the change of update time for four sensors nodes learning over hundreds of episodes. We define an episode as one transmission cycle from the device, i.e., the device transmits an update, receives a command, and enters sleep mode. In the first 200 episodes, the mechanism quickly increases IoT devices' update time because the accuracy of the information collected is below the selected margin. After the first 200 episodes, at which point the accuracy of the information collected is close to the chosen target, the increase becomes more gradual as the mechanism decides to keep the update rate constant, with small fluctuations.

Next, we analyse average energy savings an IoT device can achieve. We obtain energy consumption parameters from [17], where authors calculated energy consumption for devices using LoRaWAN. We determine the energy efficiency by comparing energy consumption between average update times obtained with the implemented updating mechanism to the baseline update time. We defined the baseline update time as the update time device requires to achieve the average accuracy of temperature measurements of ± 0.25 °C. For a fair comparison, the updating mechanism keeps average accuracy of temperature estimations within 0.25 °C of the real temperature. Figure 6 shows that with the increase in device density the energy efficiency also increases.

The proposed updating mechanism significantly improves the energy efficiency of IoT devices. Denser device deployments lead to a higher correlation exhibited in the collected data and an increase in energy efficiency.

VI. Future Work and Conclusion

The main focus of our future work will be in designing an energy-aware updating mechanism. In particular, we will focus on a scenario in which sensors rely on various energy sources, e.g., mains powered, energy harvesting etc. In such a case, the updating mechanism will have to adapt to ever-changing energy levels on devices. For example, a high frequency of updates significantly shortens a battery-powered device's lifespan, while a device connected to the power grid can operate indefinitely regardless of its frequency of updates, and the energy on the device using energy harvesting depends on external, often unpredictable, events. One possible solution is to design a reward function that will take into account the difference between the energy available at one source and the overall energy available in the system. We will strive to achieve a balanced updating mechanism, in which devices with more available energy transmit more often.

Evaluating the updating mechanism's performance using data coming from real sensing devices is crucial in our work. The SmartSantander dataset [14], used in the previous section, offers many interesting types of measurements to explore, e.g., noise pollution, light, etc. However, the used dataset is limited due to a relatively small number of sensors. Examining the dataset provided by the Intel Berkeley Research laboratory [18] enables us evaluating the updating mechanism in a network with more sensors. Another dataset we plan to adopt for the evaluation process is data collected by the Pervasive Nation[1] IoT testbed in Ireland. The higher the variety and quantity of available data, the more relevant our evaluation will be.

In this paper, we proposed an updating mechanism capable of improving energy efficiency by reducing the frequency with which the sensing devices transmit their updates. The proposed mechanism takes advantage of correlation, extracted from information collected, to improve the energy efficiency without lowering the accuracy of available information. We evaluated the proposed mechanism using data collected in SmartSantander and demonstrated that every device, collecting correlated information, can prolong its lifetime.

Acknowledgements

This work was funded in part by the European Regional Development Fund through the SFI Research Centres Programme under Grant No. 13/RC/2077 SFI CONNECT and by the SFI-NSFC Partnership Programme Grant Number 17/NSFC/5224.

References

[1] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of things for smart cities," *IEEE Internet of Things journal*, vol. 1, no. 1, pp. 22–32, 2014.

[2] L. Da Xu, W. He, and S. Li, "Internet of things in industries: A survey," *IEEE Transactions on industrial informatics*, vol. 10, no. 4, pp. 2233–2243, 2014.

[3] U. Raza, P. Kulkarni, and M. Sooriyabandara, "Low power wide area networks: An overview," *IEEE Communications Surveys & Tutorials*, 2017.

[4] "Sx1272/73 - 860 mhz to 1020 mhz low power long range transceiver datasheet," March 2015, revision 3. [Online]. Available: https://www.semtech.com/uploads/documents/sx1272.pdf

[5] R. S. Sutton and A. G. Barto, *Introduction to reinforcement learning*. MIT press Cambridge, 1998.

[6] G. Anastasi, M. Conti, M. Di Francesco, and A. Passarella, "Energy conservation in wireless sensor networks: A survey," *Ad hoc networks*, vol. 7, no. 3, pp. 537–568, 2009.

[7] T. Rault, A. Bouabdallah, and Y. Challal, "Energy efficiency in wireless sensor networks: A top-down survey," *Computer Networks*, vol. 67, pp. 104–122, 2014.

[8] D. Chu, A. Deshpande, J. M. Hellerstein, and W. Hong, "Approximate data collection in sensor networks using probabilistic models," in *ICDE*. IEEE, 2006, pp. 48–48.

[9] A. Deshpande, C. Guestrin, S. R. Madden, J. M. Hellerstein, and W. Hong, "Model-driven data acquisition in sensor networks." VLDB, 2004, pp. 588–599.

[10] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing age of information in vehicular networks," in *SECON*. IEEE, 2011.

[11] J. Hribar, M. Costa, N. Kaminski, and L. A. Da Silva, "Using correlated information to extend device lifetime," *IEEE Internet of Things Journal*, 2018.

[12] N. Cressie and H.-C. Huang, "Classes of nonseparable, spatio-temporal stationary covariance functions," *Journal of the American Statistical Association*, vol. 94, no. 448, pp. 1330–1339, 1999.

[13] T. Gneiting, "Nonseparable, stationary covariance functions for space–time data," *Journal of the American Statistical Association*, vol. 97, no. 458, pp. 590–600, 2002.

[14] V. Gutiérrez et al., "Co-creating the cities of the future," *Sensors*, vol. 16, no. 11, p. 1971, 2016.

[15] V. Mnih *et al.*, "Playing Atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.

[16] J. Hribar, A. Marinescu, N. George A, and L. A. DaSilva, "Using deep Q-learning to prolong the lifetime of correlated internet of things devices," *arXiv preprint:1902.02850*, 2019.

[17] M. Costa, T. Farrell, and L. Doyle, "On energy efficiency and lifetime in low power wide area network for the internet of things," in *IEEE, CSCN 2017*.

[18] P. Bodik, W. Hong, C. Guestrin, S. Madden, M. Paskin, and R. Thibaux, "Intel lab data," *Online dataset*, 2004.

[1] For more information; the reader may visit: https://connectcentre.ie/pervasive-nation