EmbientLoRa: Embedded Intelligence for Predictive Energy Harvesting and Management in LoRa Networks

Sukanya Jewsakul
The University of Hong Kong
sukanya@connect.hku.hk

Edith C. H. Ngai
The University of Hong Kong
chngai@eee.hku.hk

Abstract

The prolonged lifetime of energy-harvesting (EH) LoRa networks requires that all EH LoRa sensors make the best use of available harvested energy in an energy-neutral manner to avoid power failures. This requirement is challenging to be fulfilled due to the unpredictability of ambient energy sources and the spatio-temporal heterogeneity of sensors' harvesting abilities. We present EmbientLoRa, a novel predictive energy-management framework that enables energy-neutral operation in EH LoRa networks via embedded intelligence. To achieve highly-accurate EH predictions, EmbientLoRa adopts a simple yet effective technique that allows each sensor to implement a machine learning pipeline locally at low cost. Coupled with adaptive EH management, it allows the sensors with higher harvested energy to transmit critical data more frequently in a probabilistic manner without sacrificing their lifetimes. Compared with the state-of-the-art, the results from testbed experiments reveal that EmbientLoRa improves EH prediction accuracy and transmission overhead, both by up to 1.5 times.

1 Introduction

Ambient energy such as solar and radio frequency waves has been perceived as a perpetual power source [13] to support self-sustained Internet-of-Things (IoT). Nevertheless, the sporadic availability of ambient energy sources [3] has made resource-constrained EH sensors suffer from power failures. Besides the loss of data and system states, the power failures cause a *forced* shutdown and a system reset, which lead to so-called *intermittent operation* [11] that reduces the reliability of IoT systems [7]. System reset is undesirable in EH LoRa networks that implement LoRa Wide Area Network (LoRaWAN) protocol because it triggers the over-the-air activation mechanism to regain a network connection.

To avoid power failures (i.e., to prolong the lifetime of EH LoRa networks), energy-management algorithms can be employed, so as to match sensors' energy consumption (e.g., the occurrence of data transmissions) to the harvested-energy availability. Ideally, EH LoRa sensors should manage their energy consumption as close as the maximum ambient energy available in the environment to fully utilize the ambient energy resources. Unlike traditional battery-powered LoRa networks, the lifetime of EH LoRa networks can be prolonged as long as all EH LoRa sensors maintain *energy-neutral operation* (i.e, the energy consumption is less than the harvested energy). Although the existing energy-efficient techniques [8, 9, 10] for battery-powered LoRa networks can help EH LoRa sensors consume the harvested energy more efficiently, they do not necessarily optimize the ambient energy utilization and/or the application performance.

Indeed, achieving *energy neutrality* and high ambient-energy utilization is challenging as ambient energy is dynamic and varies across time and space [11]. Despite being co-located, the ambient energy concurrently harvested by the EH sensors can be completely different [3]. To cope with this challenge, sophisticated EH prediction [17, 18] and management [6, 7, 15] algorithms were proposed. Among those, the machine learning (ML)-based approaches manifest a great potential as they can predict the future harvested-energy availability for adaptive energy management. However, as the embedded EH sensors are unable to learn the ML models locally from scratch [5], the on-site data collection for offline model training, the delay of adjusting the pre-trained models to local EH conditions, and/or the communication of EH information to the (edge) server for model update/training are unavoidable for enabling accurate EH predictions. To relieve such overheads, a ML pipeline should be implemented directly on the sensors to minimize human interventions and extra communication costs.

We propose EmbientLoRa, a novel predictive energy-management framework that enables energy-neutral operation in **Em**bedded am**bient**-powered **LoRa** networks via embedded ML. It prolongs the lifetime of EH LoRa networks by ensuring that all EH LoRa sensors spend available ambient energy on useful data transmissions in an energy-neutral manner. Precisely, EmbientLoRa takes advantage of the spatio-temporal heterogeneity of ambient energy faced by individual sensors to control over data transmissions. Given the sensors harvest energy differently, those with abundant (future) harvested energy can transmit data in a probabilistic

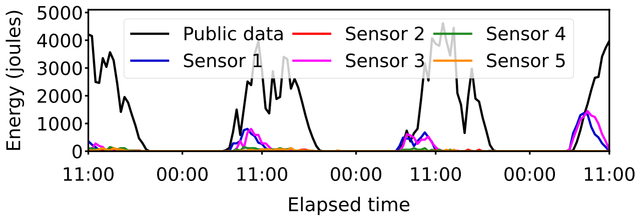

Figure 1. Heterogeneity of harvested energy

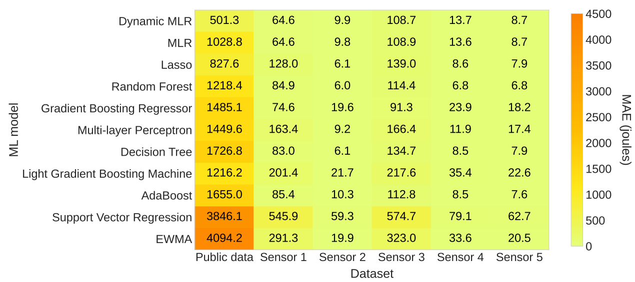

Figure 2. MAE of different ML models

manner with higher transmission probability. The transmission probability is dynamic. It is based on EH predictions from an embedded ML model on each sensor. Under low EH conditions, the sensors can only transmit critical data with low transmission probability to avoid power failures.

For the sake of dynamic energy management, EmbientLoRa enables embedded ML for EH predictions on EH LoRa sensors. Through the on-device collection of tiny EH dataset and the dynamic selection of explanatory variables (aka features or lags), each sensor can train a multiple linear regression (MLR) model for highly-accurate EH predictions locally from scratch using gradient decent (GD). Whenever a new EH sample becomes available, the model is then updated online using stochastic gradient decent (SGD) to rapidly adapt to changing EH conditions. At any time, the sensor can train a new MLR model if the prediction accuracy of the current model drops significantly. Training the new model is beneficial when the current lags are no longer a good predictor under current EH conditions. Validated through practical experiments, the experimental results show that EmbientLoRa improves EH prediction accuracy and transmission overhead, both by up to 1.5 times, compared with the state-of-the-art algorithms.

In summary, we make the following contributions:

- We present EmbientLoRa, an energy-neutral transmission management framework that prolongs the lifetime of EH LoRa networks via embedded ML. It allows EH LoRa sensors to transmit critical data in a probabilistic manner according to predictive EH availability.
- We are the first to overcome the spatio-temporal heterogeneity of ambient energy by enabling adaptive feature selection with ML model training and update for EH predictions fully on embedded EH LoRa sensors. The highly-accurate EH predictions lead to an improvement in ambient energy utilization and management.

2 Motivation

Although existing work on embedded ML for EH predictions has allowed EH sensors to update model weights [5, 18] or prune inaccurate trees from random forests [17] locally, on-device feature selection and model training still remain an open challenge. The dynamic selection of features is essential to identify relevant input variables for accurate EH predictions. For example, in solar EH systems, the solar energy harvested at any given time of the day can be estimated from the solar energy harvested at the same time on previous days [1]. This estimation, namely exponentially weighted moving average (EWMA), is appropriate when daily EH conditions are stable. However, under varying EH conditions, the solar energy harvested at different times on the same day should also be considered. Given solar energy varies across time and space, each EH sensor should be able to select relevant features and train a new ML model locally whenever the EH-prediction accuracy of the current model significantly drops.

As an example of empirical EH measurements, Figure 1 shows the spatio-temporal heterogeneity of solar energy concurrently harvested by five EH sensors which were deployed on windowsills in the same building for 72 hours. To differentiate between our own solar-energy measurements and solar energy estimated from the public data [14] of *global solar radiation* measured from a station in our region, we also calculated the solar energy that our EH systems could generate based on the solar radiation available at the same times on those days. Noticeably, the solar energy calculated from public data differs from our measurements. Without (long-term) EH data collected from deployment locations, the efficiency of existing work [17, 18] that relied on public EH data for offline feature selection and model training could consequently decline. To achieve highly-accurate EH predictions, solar energy measured from actual deployment locations should thus be included during feature selection and model training.

To motivate the need for our work, Figure 2 compares the prediction performance of different ML algorithms. Each algorithm was trained individually using the corresponding time-series of solar energy shown in Figure 1. Unlike other algorithms that were trained offline using PyCaret [2], Dynamic MLR (aka D-MLR algorithm used in our proposed framework), MLR, and EWMA [5] algorithms were handcrafted using Python and supported online learning. Unlike other algorithms where relevant features were decided once, D-MLR algorithm identified features and trained a new MLR model dynamically whenever the accuracy of EH predictions significantly dropped. Here, the autocorrelation function (ACF) was used to determine relevant features from EH data. Compared with other algorithms, D-MLR algorithm offers the lowest mean absolute error (MAE) values in most cases when predicting the solar energy generated in the next hour over a few days.

3 Related Work

EH prediction algorithms provide insights into future ambient-energy availability which are beneficial for careful energy management. With recent advances in embedded ML, adaptive EH prediction algorithms have been pro-

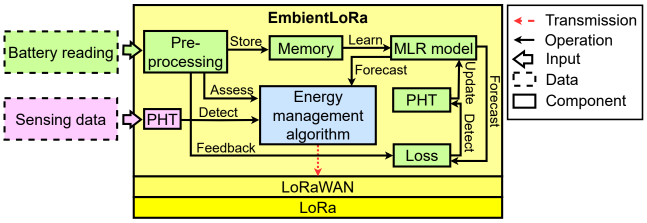

Figure 3. Operational architecture of EmbientLoRa

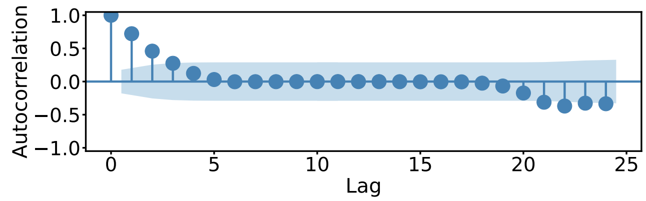

Figure 4. Temporal autocorrelation of harvested energy

posed to cope with the spatio-temporal heterogeneity of ambient energy. For example, OB-HEP [5] enhanced the accuracy of EH predictions by learning EWMA weight for each timeslot locally on an EH sensor using gradient descent. In [18], a hierarchical EH prediction algorithm that combined eight pre-trained neural networks for energy-level and energy-availability predictions was proposed. To gradually adapt to local EH conditions, the neural network weights were updated using a backpropagation algorithm when small EH datasets were available. Moreover, [17] adopted the weighted average of the output from immutable random forests' trees for indoor EH predictions. By updating the weights of the trees online, inaccurate trees could be pruned from random forests. Aiming to prolong the lifetime of an EH LoRa network, a medium access control (MAC) protocol called Long-Lived LoRa (LLL) [6], on the other hand, allowed EH LoRa sensors with superfluous energy to relay packets from neighboring sensors with depleting energy to LoRa gateways. By reserving low spreading factors (SFs) for packet offloading and dividing the network into smaller cells with a fixed channel each, LLL helped depleting sensors offload packets to superfluous sensors at low cost and (inter-cell) interference.

Different from the above work, EmbientLoRa implements dynamic feature selection and ML model training/update for accurate EH predictions locally on embedded EH LoRa sensors. Through energy-neutral transmission management, it prolongs the lifetime of individual sensors by having them to transmit critical data in a probabilistic manner in response to EH predictions from their embedded MLR models.

4 EmbientLoRa

4.1 System Architecture

Figure 3 illustrates the operational architecture of EmbientLoRa, a predictive energy-management framework that is designed to enable energy-neutral operation in embedded EH LoRa networks. EmbientLoRa is composed of two components: embedded intelligence for accurate EH predictions and adaptive transmission management. Being an application-layer framework, it is compatible with an off-the-shelf LoRa sensor powered by ambient energy.

For the sake of simplicity, we consider that each EmbientLoRa sensor is equipped with a solar EH system (i.e., a solar panel and a rechargeable battery). To avoid unstable power supply, the EH LoRa sensor is powered by the rechargeable battery where the harvested solar energy is stored and estimated. At the time of deployment, the EH system specification (e.g., solar-panel area, the size of rechargeable battery, and voltage/current at the maximum power) is made known to each sensor for the purpose of data transformations.

4.2 Energy Prediction

Data Preparation. To train a multiple linear regression (MLR) model (aka D-MLR model in Section 5) for EH predictions locally on an EH LoRa sensor, we divide time of day into 24 timeslots of one hour each. With that, we program the sensor to measure the hourly harvested energy for each timeslot and update a local dataset with a new sample regularly. The solar energy harvested during the past hour can be estimated from battery readings using Equation 1.

$$E_t^h = E_t^b - (E_{t-1}^b - E_{t-1}^c)\,, where\, E_t^h \geq 0. \tag{1}$$

Specifically, the EmbientLoRa sensor wakes up regularly every transmission cycle (T_{cycle}) to perform tasks related to main application and energy management. At the beginning of each T_{cycle}, after collecting sensing data, the sensor measures its battery capacity and estimates the energy harvested in the past T_{cycle} duration. At any time t, the harvested energy (E_t^h) can be estimated by subtracting the difference between the battery capacity measured at the beginning of previous transmission cycle (E_{t-1}^b) and the total energy consumed in that cycle (E_{t-1}^c) from the current battery measurement (E_t^b). To estimate E_{t-1}^c, the sensor can employ the power profile of different system states (i.e, transmitting, receiving, and sleep states) provided at the time of deployment. If T_{cycle} is shorter than an hour, E_t^h is accumulated and hourly stored in a dataset for model training or update. Before being stored, the hourly harvested energy could be transformed (i.e., preprocessed) in the form of a percentage in proportion to the maximum amount of energy the designed solar EH system can generate within an hour under a full-sunlight condition, so as to enhance the accuracy of EH predictions. During the transformation, the EH system specification is leveraged.

On-device Model Training. To trade off computational overhead against prediction accuracy, we allow the sensor to keep the history of hourly harvested energy for the last three days (i.e., 72 hours) in its memory. Maintaining tiny dataset helps accelerate ML model training while minimizing memory usage. We derive at this number of days by analysing both our own solar EH trace collected from an EH sensor over six days and two public datasets that cover two month-long EH data [14] in our region and a year-long EH data [12] from National Renewable Energy Laboratory (NREL). For example, as shown in Figure 4, the ACF of hourly harvested energy from our own empirical measurements becomes less relevant after two days passed. Aligning with the previous study [18], we do not notice any significant improvement in prediction accuracy when testing with larger datasets.

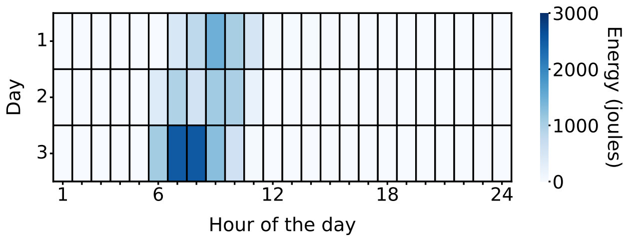

Figure 5. Heat map of self-collected harvested energy

Algorithm 1: Transmission management

```
Input: D_t, E_t^b, E_t^min, S_h, S_tx
TX ← False
if E_t^b ≥ E_t^min then
    if S_h ≥ 1 or CRITICAL(D_t) is True then
        if RANDOM() ≤ S_tx then
            TX ← True
        end
    end
end
if TX is True then
    Transmit D_t
else
    Abort transmitting D_t
end
```

Whenever the EH dataset is ready (see Figure 5), the MLR model can be trained on the EH LoRa sensor from scratch. Precisely, the sensor uses ACF to detect significant features including seasonal and non-seasonal lags. The seasonal lags concern past observations at the same timeslot on previous days. The non-seasonal lags, on the other hand, involve past observations at previous timeslots on the same day. Due to the size of EH dataset, the maximum number of seasonal lags are limited to two. In contrast, up to three non-seasonal lags can be chosen. As shown in Figure 4, EH data become less relevant after three consecutive hours passed. Consequently, up to five lags (i.e., two seasonal lags and three non-seasonal lags) can be used for ML-model training. After significant lags are determined, the MLR model is trained using gradient decent where ordinary least squares method is used as a loss function. Similar to [5], all regression coefficients are initialized to 0.5 when a training process starts. The learning rate and the number of iterations, on the other hand, are set to 0.3 and 20, respectively. We allow the training process to be early terminated if all gradients are less than 0.001. Once the training process is done, the MLR model is deployed for EH predictions and management. To cope with power failures, the trained MLR model and the EH dataset can be stored in non-volatile memory regularly. After recovering from a power failure, the sensor can consequently load the stored model and EH data from the memory for EH predictions. By progressively learning from new samples, the accuracy of EH predictions will be improved over time.

Online Learning. To keep the embedded MLR model up-to-date, the sensor estimates a prediction error (i.e., loss) regularly whenever the actual hourly harvested energy becomes available. With that, the two-sided Page-Hinkley test (PHT) algorithm [16] is employed to diagnose a drift in the mean of losses. In case the concept drift is detected (i.e., the mean of losses changes) and the EH dataset is ready, the sensor uses ACF to find a new set of features and trains a new MLR model accordingly using gradient descent. In an absence of concept drift, the current MLR model is updated online from a new sample using stochastic gradient descent. Relying on a weighted sum method, we use the forgetting rate which is set to 0.3 as a weighting factor to gradually forget past coefficients and adapt to current EH conditions.

Inference. The MLR model employs a multi-step prediction approach to predict the total amount of energy potentially harvested from any time t to the end of day (denoted as T; $t \leq T$). Namely, the prediction period is in the range [t, T]. We denote the predicted amount of energy harvested during this time period as E_t^p, where t symbolizes the starting time of multi-step predictions. In the first prediction step, if the starting time t is somewhere behind the beginning of an hour-long timeslot (i.e., the prediction duration of the first timeslot is shorter than an hour), the amount of energy predicted for the first timeslot is linearly interpolated using the hourly harvested energy predicted for that timeslot. Then, the predicted EH availability (i.e., E_t^p) is used by an energy-management algorithm to make a transmission decision.

4.3 Energy Management

Algorithm 1 outlines our proposed energy-neutral transmission management. To decide whether or not to transmit the sensing data (D_t), the EmbientLoRa sensor calculates a transmission score (S_{tx}) in the range [0, 1]. With that, it performs a data transmission in a probabilistic manner according to S_{tx} score calculated (see Equation 5). Fundamentally, S_{tx} score is the normalization of harvesting score (S_h). The S_h score rates how well the total EH availability (E_t^a) can accommodate the total energy consumption (E_t^e).

Besides the predicted EH availability (i.e., E_t^p discussed above), the sensor also keeps track of the actual harvested energy (E_t^r) residing in the energy storage to assess the total harvested energy (E_t^a) available from any time t to T, where T denotes the end of day. According to Equation 2, E_t^r is basically the accumulated difference between battery capacities measured at the current and the past T_{cycle} intervals.

$$E_t^a = E_t^p + E_t^r \text{, } where\, E_t^r = max(0, E_{t-1}^r + E_t^b - E_{t-1}^b) \quad (2)$$

The total energy consumption (E_t^e), on the other hand, is a product of the energy consumed within each transmission cycle (E_{cycle}) and the number of remaining transmission cycles on that day (see Equation 3). Basically, E_{cycle} involves the energy consumed by a single data transmission with two subsequent receive windows and a low-power sleep for the rest of T_{cycle} duration.

$$E_t^e = E_{cycle}(\left\lceil \frac{T-t}{T_{cycle}} \right\rceil + 1) \quad (3)$$

After E_t^a and E_t^e are estimated, the sensor can calculate the harvesting score (S_h) using Equation 4. Basically, S_h score is the ratio of E_t^a to E_t^e. By definition, energy neutrality should be achieved if S_h score is greater than one (i.e., the energy consumption is smaller than the harvested energy).

$$S_h = \frac{E_t^a}{E_t^e} \quad (4)$$

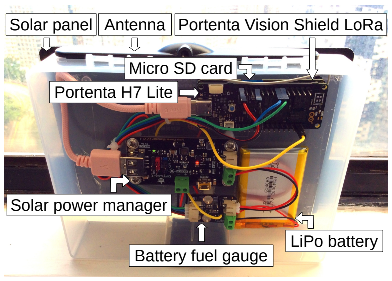

Figure 6. Testbed device and deployment

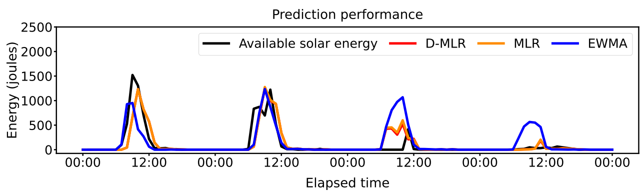

Figure 7. Prediction accuracy of different ML models

Finally, the transmission score (S_{tx}) can be calculated from S_h score. As shown in Equation 5, whenever S_h score increases, S_{tx} score increases accordingly. Consequently, the sensor having high EH availability can transmit (critical) sensing data with high transmission probability. We use the two-sided PHT algorithm to detect the criticality of sensing data (D_t). To avoid power failures, no transmissions occur if the battery capacity (E_t^b) at time t does not meet a user-defined minimum backup-energy requirement (E_t^{min}).

$$S_{tx} = 1 - \frac{1}{1+S_h} \tag{5}$$

5 Testbed Experiments and Results

5.1 Implementation

We realized EmbientLoRa on off-the-shelf LoRa sensors powered by solar EH systems (see Figure 6). As our EH LoRa sensors, we used Arduino Portenta H7 Lite boards. Each board is characterized by STM32H747 32-bit dual core processor running at 480 MHz (which was set to run at 60 MHz) and 240 MHz (which was disabled), respectively. This board supports MicroPython programming which helps facilitate the implementation of our own embedded ML (i.e., D-MLR) libraries. Additionally, MLR and EWMA [5] libraries (see Section 2 and 3) were also implemented, so as to further compare the performance of different embedded ML models. To enable LoRa communications, add-on boards called Arduino Portenta Vision Shield LoRa were also employed. Each add-on board contained a micro SD card which we used to store system log and dataset for ML model initialization. Being self-powered, each sensor was connected to a DF Robot Solar Power Manager 5V board via USB. The power management board was then connected to a 6V 0.5A solar panel, a 3.7V 1200mAh LiPo battery, and a Gravity 3.7V Lithium Battery Fuel Gauge module used for battery measurements. We used 8-channel RAK7268 WisGate Edge Lite 2 as a LoRa gateway, and deployed ChirpStack LoRaWAN network server on our personal computer.

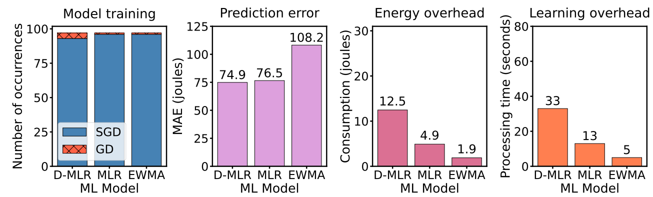

Figure 8. Performance of different ML models

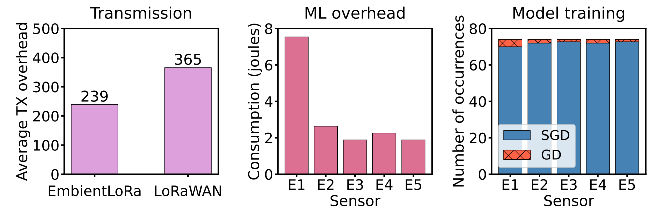

Figure 9. Transmission overhead (left). ML overhead (middle). Training occurrences (right)

5.2 System Setup and Deployment

As a proof of concept, we first equipped an EH LoRa sensor with embedded ML models discussed above. We then used eight-day-long solar EH measurements from Sensor 2 in Figure 1 to practically compare the performance of those models. Additionally, we also deployed five pairs of **L**oRaWAN [4] and **E**mbientLoRa (equipped with D-MLR model) sensors across three floors in our departmental building to collect empirical data for three days. As the solar panels of our choice could not harvest energy from light bulbs, we concurrently placed each pair of EmbientLoRa and LoRaWAN sensors (e.g., E1 and L1) next to each other on the same windowsill, so as to allow them to be exposed to comparable sunlight conditions. To speed up the experiment, we used the tiny EH dataset shown in Figure 5 to initialize a ML model for all EmbientLoRa sensors. Consequently, the sensors could immediately learn the ML models from their local EH conditions after the deployment.

5.3 Experimental Results

5.3.1 *ML Performance*

Figure 7 validates the EH prediction performance of our D-MLR model against MLR and EWMA [5] models. Compared with EWMA and MLR models, our D-MLR model offers higher prediction accuracy (see also the second diagram of Figure 8), particularly when the daily EH conditions changed significantly during the third day of EH predictions. Given the scarce solar-energy availability on that day, the highly-accurate EH predictions from D-MLR model could help an EH LoRa sensor carefully manage both incoming and existing energy (i.e., the future harvested energy and the residual energy) to avoid power failures. Although the dynamic feature selection and model training using GD (see the first diagram of Figure 8) allow D-MLR model to rapidly adapt to abrupt changes in EH conditions, an improvement in EH prediction accuracy comes at the expense of slightly higher computational overhead. As illustrated in the last two diagrams of Figure 8, the processing overhead and the energy consumption of D-MLR model over the course of eight

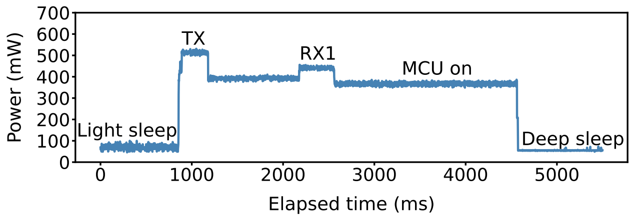

Figure 10. Power consumption profile of testbed device

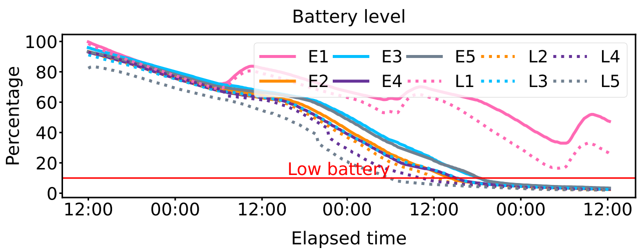

Figure 11. Time-series of battery levels

days could be slightly higher than MLR and EWMA models.

5.3.2 *Network Performance and Lifetime*

We define transmission overhead as the number of data (re)transmissions which were averaged per sensor. As can be seen from the leftmost part of Figure 9, due to probabilistic-based transmissions, EmbientLoRa offers significantly lower transmission overhead. Precisely, EmbientLoRa nearly cuts the transmission overhead of LoRaWAN by half. Although EmbientLoRa sensors performed feature selections and trained ML models locally, the cost of embedded ML (i.e., processing) could be negligible when being compared with the cost of data transmissions (see Figure 10). Precisely, Figure 9 shows the energy consumed by on-device model training. As can be seen from the middle and rightmost parts of Figure 9, EmbientLoRa sensors consumed small amount of energy for ML. Essentially, EmbientLoRa sensors performing on-device feature selections and model training using GD more frequently exhibit the higher energy consumption. Compared with LoRaWAN sensors, EmbientLoRa sensors have longer lifetimes. As shown in Figure 11, a few EmbientLoRa sensors reached the low-battery threshold, which was set to 10 percent, after two days of the deployment passed. Due to probabilistic-based transmissions, EmbientLoRa sensors transmitted data less frequently. Consequently, despite being deployed in similar EH environments, LoRaWAN sensors depleted their batteries faster than EmbientLoRa sensors.

6 Conclusion

This paper presents a novel energy-neutral transmission management framework called EmbientLoRa. It prolongs the lifetime of EH LoRa networks via embedded ML. Through the on-device selection of relevant features, EmbientLoRa lets each sensor to dynamically train a ML model for highly-accurate EH predictions, which lead to an improvement in ambient energy utilization and management. By examining the predictive EH availability, EmbientLoRa allows the sensors with higher EH availability to transmit critical data more frequently in a probabilistic manner. The results from testbed experiments reveal that EmbientLoRa outperforms all the existing methods.

7 References

[1] K. S. Adu-Manu, N. Adam, C. Tapparello, H. Ayatollahi, and W. Heinzelman. Energy-harvesting wireless sensor networks (eh-wsns): A review. *ACM Trans. Sen. Netw.*, 14(2), apr 2018.

[2] M. Ali. Pycaret: An open source, low-code machine learning library in python, 2020.

[3] A. Cammarano, C. Petrioli, and D. Spenza. Online energy harvesting prediction in environmentally powered wireless sensor networks. *IEEE Sensors Journal*, 16(17):6793–6804, 2016.

[4] L. A. T. Committee. Lorawan® l2 1.0.4 specification (ts001-1.0.4). *LoRa Alliance: Fremont, CA, USA*, 2020.

[5] K. Dolui, A. S. Thangarajan, T. Claes, S. Michiels, and D. Hughes. Towards on-board learning for harvested energy prediction. In *Proceedings of the 6th International Workshop on Embedded and Mobile Deep Learning*, EMDL '22, pages 7–12, New York, NY, USA, 2022. Association for Computing Machinery.

[6] S. Fahmida, V. P. Modekurthy, M. Rahman, A. Saifullah, and M. Brocanelli. Long-lived lora: Prolonging the lifetime of a lora network. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*, pages 1–12, 2020.

[7] F. Fraternali, B. Balaji, D. Sengupta, D. Hong, and R. K. Gupta. Ember: Energy management of batteryless event detection sensors with deep reinforcement learning. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, SenSys '20, pages 503–516, New York, NY, USA, 2020. Association for Computing Machinery.

[8] A. Gamage, J. C. Liando, C. Gu, R. Tan, and M. Li. Lmac: Efficient carrier-sense multiple access for lora. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, MobiCom '20, New York, NY, USA, 2020. Association for Computing Machinery.

[9] W. Gao, W. Du, Z. Zhao, G. Min, and M. Singhal. Towards energy-fairness in lora networks. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pages 788–798. IEEE, 2019.

[10] W. Gao, Z. Zhao, and G. Min. Adaplora: Resource adaptation for maximizing network lifetime in lora networks. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*, pages 1–11, 2020.

[11] K. Geissdoerfer and M. Zimmerling. Learning to communicate effectively between battery-free devices. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 419–435, Renton, WA, Apr. 2022. USENIX Association.

[12] N. R. E. Laboratory. Nrel national solar radiation database, 2000.

[13] A. Y. Majid, P. Schilder, and K. Langendoen. Continuous sensing on intermittent power. In *2020 19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 181–192, 2020.

[14] H. K. Observatory. Regional weather in hong kong — the latest 1–minute global solar radiation and direct solar radiation and diffuse radiation, 2023.

[15] N. Saoda, W. Wang, M. F. R. M. Billah, and B. Campbell. An energy supervisor architecture for energy-harvesting applications. In *2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 323–336, 2022.

[16] R. Sebastião and J. M. Fernandes. Supporting the page-hinkley test with empirical mode decomposition for change detection. In M. Kryszkiewicz, A. Appice, D. Ślezak, H. Rybinski, A. Skowron, and Z. W. Raś, editors, *Foundations of Intelligent Systems*, pages 492–498, Cham, 2017. Springer International Publishing.

[17] N. Stricker and L. Thiele. Accurate onboard predictions for indoor energy harvesting using random forests. In *2022 11th Mediterranean Conference on Embedded Computing (MECO)*, pages 1–6, 2022.

[18] N. Yamin and G. Bhat. Online solar energy prediction for energy-harvesting internet of things devices. In *2021 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*, pages 1–6, 2021.