Analysing and Improving Robustness of Predictive Energy Harvesting Systems

Conference Paper**Author(s):**

Stricker, Naomi Thiele, Lothar

Publication date:

2020-11

Permanent link:

<https://doi.org/10.3929/ethz-b-000457332>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

<https://doi.org/10.1145/3417308.3430264>

Analysing and Improving Robustness of Predictive Energy Harvesting Systems

Naomi Stricker
ETH Zurich
Zurich, Switzerland
nstricker@ethz.ch

Lothar Thiele
ETH Zurich
Zurich, Switzerland
thiele@ethz.ch

ABSTRACT

Internet of Things (IoT) systems can rely on energy harvesting to extend battery lifetimes or even to render batteries obsolete. Such systems employ an energy scheduler to optimize their behaviour and thus performance by adapting the node operation. Predictive models of harvesting sources, which are inherently non-deterministic and consequently challenging to predict, are often necessary for the scheduler to optimize performance. Therefore the accuracy of the predictive model inevitably impacts the scheduler and system performance. This fact has been largely overlooked in the vast amount of available results on energy management systems. We define a novel robustness metric for energy-harvesting systems that describes the effect prediction errors have on the system performance. Furthermore, we show that if a scheduler is optimal when predictions are accurate, it is not very robust. Thus there is a trade-off between robustness and performance. We propose a prediction scaling method to improve a system's robustness and demonstrate the results using energy harvesting data sets from both outdoor and indoor scenarios. The method improves a non-robust system's performance by up to 75 times in a real-world setting.

CCS CONCEPTS

• **Computer systems organization** → **Sensor networks; Sensors and actuators**; • **Hardware** → **Sensor applications and deployments**.

KEYWORDS

Energy Harvesting, Scheduling, Robustness, Internet of Things

ACM Reference Format:

Naomi Stricker and Lothar Thiele. 2020. Analysing and Improving Robustness of Predictive Energy Harvesting Systems. In *ENSys '20: ACM SIGBED International Conference on Embedded Software, November 16–19, 2020, Yokohama, Japan*. ACM, New York, NY, USA, 7 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

Energy harvesting is seen as a viable option to make the growing Internet of Things (IoT) more scalable by enabling long-term self-sustainable, maintenance-free operation. Nodes of such systems typically contain one or more harvesting sources, a rechargeable

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ENSys '20, November 16–19, 2020, Yokohama, Japan

© 2020 Copyright held by the owner/author(s).

<https://doi.org/0000001.0000001>

energy storage, an energy management system and the application. The energy management system controls the energy flow between the harvesting sources, the energy storage and application.

A first major challenge is the efficient use of available resources. A large number of energy schedulers have been proposed to orchestrate the use of available energy. They control the application's operation by adapting its sensing, actuation, computation and communication rate, data processing algorithms or by enabling an *energy scarcity* operation mode [1–7]. They often exploit predictions of future harvested energy to optimize long-term system utility.

A second major challenge arises from the high temporal and spatial variability that energy sources may experience. This complicates system optimization, which depends on the interplay between the harvesting environment, energy storage capacity, state-of-charge, and requested node operation. In addition, when the system is deployed to multiple locations, it will invariably face situations where previous design time assumptions do not hold. In such cases, the predictive model fails to correctly estimate the harvested energy, impacting the scheduler and the system behavior. To mitigate these challenges, variability can be absorbed by overdimensioning system resources but this violates the first challenge.

To illustrate the severe influence of prediction errors, we provide a simple example. In [6, 8], a scheduler is proposed that maximizes the minimal energy the application uses and optimizes long-term system utility. The solar harvester is placed in an outdoor environment for two years. The exponentially weighted moving average (EWMA) prediction [1] in Figure 1(a) closely follows the actual harvested energy. In fact, its mean relative prediction error is only

(a) The EWMA prediction closely matches the harvested energy. Its mean relative prediction error is only 6%.

(b) In the shaded areas, the weekly energy used by the system collapses to 66% and 73% of the perfect prediction case.

Figure 1: Even small errors in energy prediction lead to significant penalties on the weekly used energy.

6 %. Figure 1(b) shows the energy the application uses determined by the scheduler combined with the EWMA as well as a perfect prediction. The shaded regions highlight when prediction errors cause the minimum weekly used energy to significantly decrease.

Because the scheduler's performance can be sensitive, it is necessary to fine-tune the algorithm and parameters for selected environmental data sets [9]. Consequently, the resulting system configurations are typically optimal only in a limited set of possible environmental conditions. The example above highlights the need to analyse the repercussions of prediction errors.

In this work, we introduce the notion of robustness to address this gap. Robustness analyses the effect prediction errors have on the scheduler performance. It provides knowledge about the system behavior for a wider range of environmental conditions. As such, robustness provides a novel and important point of view for evaluating, designing and optimizing the system behavior. In addition, we demonstrate that a scheduler that performs optimally when predictions are accurate, does not have a high degree of robustness. Hence we identify a robustness performance trade-off. Where previously the design process relied on optimizing performance by evaluating a few selected scenarios, this approach neglects robustness and the resulting trade-off. The new dimension in the design space encourages a bi-objective optimization. We propose a prediction scaling method to explore the trade-off and improve the robustness of non-robust schedulers. The paper contains the following new results:

- We define a robustness measure that quantifies the influence prediction accuracy has on a system performance measure.
- We show that a scheduler that has optimal performance when predictions are accurate, is not robust and thus demonstrate a robustness performance trade-off.
- We propose a method that enables exploring this trade-off between robustness and efficient use of resources.
- The above results are validated by extensive simulations of several well known energy schedulers using indoor and outdoor solar harvesting data sets.

2 RELATED WORK

Predictive energy harvesting systems is a broad and rich field that has been studied for many years. Many works focus on one of the subsystems that constitute the energy management system.

Schedulers: Various algorithms have been proposed to optimize system performance for specific optimization objectives. Possible optimization objectives include deviations from energy neutral operation e.g. [5, 10], quality of service [4, 11] and sampling rate [12]. Some schedulers provide an optimal solution for their objective and system assumptions. In [8] the minimal utilized energy is maximized and optimality is achieved for perfect prediction. [4] proposes a dynamic programming algorithm that iteratively determines the optimal task schedule to optimize the summed quality the tasks provide. On the other hand, schedulers can also be based on heuristics e.g. [12], or data-driven [5]. These works on schedulers, however, consider only specific environments to determine and evaluate performance. Therefore it is not known how these algorithms perform in new environments and with diverse prediction errors.

Predictors: Energy predictors for solar energy include Expo-

nentially Weighted Moving-Average (EWMA) [1], Improved Pro-Energy [13], Q-learning based solar energy prediction (QL-SEP) [2] or Artificial Neural Network [14]. While these algorithms can exhibit high accuracy, their statistical analysis is typically only valid for the evaluated data sets. As energy harvesting systems are deployed in a wide variety of environments over long periods of time, the prediction errors will experience greater variability.

Robustness: Our proposed robustness analysis is comparable to other studies in the Wireless Sensor Networks (WSN) community. Robustness of WSNs studies the effects unpredictable links or node failures have on the network, e.g. [15–17]. In [15] the authors propose a method to improve robustness and in [16] a trade-off between robustness and network life-time is optimized. However, our focus lies on when energy predictors fail to provide accurate predictions, and the effect it has on the overall system behavior. We provide new insights on the reliability and performance of predictive energy harvesting systems.

Furthermore, we propose a novel prediction scaling method to mitigate scenarios where prediction errors result in a catastrophic collapse in system performance. Our work shows that certain schedulers are inherently robust, but a scheduler that is optimal when predictions are accurate exhibits a lower degree of robustness.

3 ROBUSTNESS ANALYSIS OF HARVESTING SYSTEMS

A general harvesting system model is described and subsequently, the concept of robustness for these systems is introduced.

3.1 System Model

A discrete-time model is used for the system's evolution, where the time horizon of interest $[0, T)$ is discretized into intervals $t \in [0, T)$. Figure 2 shows a block diagram describing the general system. The harvester converts primary energy into electrical energy thus producing $E_{harv}(t)$ during $[t, t+1)$. The harvested energy is buffered in an energy storage with finite capacity B . Its state-of-charge at time t is denoted by $b(t)$. At every time interval t the scheduler determines, based on energy predictions E_{pred} , the energy E_{sched} that the application (APP) should use in the future. The energy the application uses is the allocated $E_{sched}(t)$ if sufficient energy is available and otherwise all the available energy is used,

$$E_{used}(t) = \min(E_{sched}(t), b(t) + E_{harv}(t))$$

Thus, the energy storage state-of-charge evolves according to

$$b(t+1) = \max\{\min\{b(t) + E_{harv}(t) - E_{used}(t), B\}, 0\}$$

Figure 2: A harvesting system with an energy management system to control the energy flow. Prediction Scaling to increase robustness is discussed in Section 5.

3.2 Defining the Robustness Metric

Robustness captures the effect prediction errors have on the scheduler and thus system performance. To define robustness, we formalize comparable prediction accuracy and performance metrics.

Predictor The energy predictor estimates the future harvested energy. Because accurate predictions are challenging, the harvested energy will differ from the prediction. This deviation is defined as

$$E_{pred}(t) = E_{harv}(t) \cdot (1 + e(t)) \quad (1)$$

where $e(t)$ is the relative error. This is not always the most appropriate accuracy measure. For example, where $E_{harv}(t) = 0$ but $E_{pred}(t) \neq 0$, the prediction error in (1) is not meaningful. In such cases, other accuracy metrics such as the mean absolute deviation percentage [18] are more expressive.

Scheduler and System Performance The system performance throughout the time horizon $[0, T]$ is usually summarized by a scalar performance metric. Because schedulers determine the application behavior, they have a great influence on the system performance. Hence, a performance metric related to the scheduler optimization objective lends itself well. Examples include system utility functions that depend on the duty cycle, e.g. [1, 9], a total utility that is the concave sum of all the used energy, $U = \sum_{t \in [0, T]} \mu(E_{used}(t))$ for a strictly concave function μ or the minimal used energy per time interval $E_{min} = \min_{t \in [0, T]} E_{used}(t)$ [8]. While we focus on the latter two sample utility functions, our analysis is also applicable to other utility functions, such as events or deadlines the system misses as well as the system's energy efficiency.

Robustness The degree of robustness is a comparison of the prediction error and system performance.

Error Model The error model consists of a systematic and random error. A systematic error occurs if the prediction systematically over- or underpredicts whereas a random error represents randomly fluctuating errors. For the robustness analysis, the error is assumed to follow a normal distribution, $e(t) \sim \mathcal{N}(\mu, \sigma)$, where μ represents the systematic and σ the random error. Thus the prediction error for each time interval is a random sample from the error distribution and the prediction is constructed to be $E_{pred}(t) = \max(E_{harv}(t) \cdot (1 + e(t)), 0)$. The analysis can be performed for other error distributions, such as skewed or heavy-tailed distributions. However, the extensive case study summarized in Section 4.5 shows that the error distribution does not have a great influence on the results and conclusions.

Performance The performance metric is determined for the prediction generated from T random samples from the error distribution. Each realization leads to a possibly different performance. The performance metric is therefore also a random variable however with an unknown distribution. If an analytical solution for the performance distribution can be derived, the subsequent robustness analysis is performed analytically. Otherwise a numerical approach is necessary. N Monte Carlo simulations are performed to generate N independent samples of the performance metric distribution. Subsequently, these are used for the robustness analysis.

Robustness The distributions of the prediction error and system performance are compared based on percentiles, e.g. the median. From the N independent samples of the performance metric distribution a confidence interval for the percentile is determined based on non-parametric statistics using the method proposed in [19].

To enable a meaningful comparison, the relative changes in the percentiles are compared. Thus the effect prediction errors from one distribution have on the performance is described by

$$r = \left| \frac{\Delta Q_{error}}{\Delta Q_{perf}} \right| \quad (2)$$

where $\Delta Q_{error} = \frac{Q_{error(0)} - Q_{error}}{Q_{error(0)}}$ is the relative change in the prediction error percentile and $\Delta Q_{perf} = \frac{Q_{perf(0)} - Q_{perf}}{Q_{perf(0)}}$ the relative change in the performance metric percentile. Q_X is the percentile of the random variable X . And $perf(0)$ denotes the performance metric for a prediction error without systematic error $error(0)$, and $perf$ is the performance metric of the analysed error distribution $error$. The smaller the effect prediction errors have on the performance, the larger r is and vice versa. Thus, the larger the metric as defined in (2) is, the more robust the system is against prediction errors from the analysed error distribution. Building on (2), the degree of robustness against a set of error distributions is defined as

$$R = \min_{PDF_{error}} \left| \frac{\Delta Q_{error}}{\Delta Q_{perf}} \right| \quad (3)$$

where ΔQ_{error} and ΔQ_{perf} are defined as for (2). In addition, PDF_{error} denotes all the error distributions in the set that is being analysed. Similar to r , as R increases the largest effect that prediction errors from the considered set have on the performance decreases and vice versa. In conclusion, the larger R is the more robust the system is against the analysed set of prediction errors.

4 CASE STUDY: ROBUSTNESS ANALYSIS

The case study investigates three schedulers to compare them with respect to robustness and performance. Further the optimal scheduler is shown to have a low degree of robustness thus highlighting a robustness performance trade-off.

4.1 Schedulers and Performance Metrics

We analyse three schedulers that adapt the application energy and have comparable objectives. For the duty cycle adaption (DCA) algorithm [1] the time interval is set to 1 hour and the minimal duty cycle to 0.01. The long-term energy neutral operation (LT-ENO) algorithm [9] has the same minimal duty cycle. Thirdly, the optimal power management with guaranteed minimal energy utilization (OPT) [8] directly determines the scheduled energy E_{sched} . The scheduler can be implemented in a resource-efficient manner using a look-up table [6, 8].

Table 1: Name, time-period, and location of the data sets. P_{max} is the maximum possible system consumption, A_{pv} is the panel size and Ω the environmental parameter.

Name	Time Period	Lat [°]	Long [°]	P_{max} [W]	A_{pv} [cm ²]	Ω
AK	1/61 – 12/72	61.21	-149.90	0.128	20	0.3448
CA	1/98 – 12/09	34.05	-117.95	0.353	10	0.3282
MD	1/61 – 12/72	39.29	-76.61	0.3915	15	0.3351
MI	1/98 – 12/09	42.05	-86.05	0.286	15	0.4729
ON	1/98 – 12/09	48.05	-87.65	0.248	20	0.4855
OR	1/61 – 12/72	45.52	-122.67	0.27	15	0.3955
TX2	1/61 – 12/72	28.05	-97.39	0.4135	10	0.3431
TX1	1/61 – 12/72	31.77	-106.48	0.3915	15	0.2242

Figure 3: The three schedulers lead to different daily used energy curves for the same harvesting trace.

The two performance metrics used are the minimal used energy per day and the total utility, $U = \sum_t \mu(E_{used}(t))$ where $\mu(x) = \sqrt{x}$ [8]. The first captures the minimal service that the system provides and the latter the overall system utility.

4.2 Setup and Data

The robustness analysis is performed using outdoor energy traces and synthetic predictions. The data set stems from the publicly available National Solar Radiation Database [20] and is summarized in Table 1. The energy storage capacity is $B = 143$ Wh and has a discharging efficiency of $\eta_{out} = 0.7$. This efficiency is incorporated by scaling the energy flowing out of the energy storage with η_{out} [21]. Other inefficiencies, for example leakage power, and energy storage hysteresis are not regarded. We analyse the robustness against a set of error distributions with systematic error $\mu \in [0, 0.2]$, and random error $\sigma = 0.05$. With errors drawn from the error distribution, the prediction is constructed to be $E_{pred}(t) = \max(E_{harv}(t) \cdot (1 + e(t)), 0)$. We compare the prediction error median, not its relative change because at zero systematic error the median is zero, to the relative change in the performance median. The latter is estimated by running a large number, $N = 100$, of Monte Carlo simulations and determining the 95 % confidence interval of the median using the Thomson method [19]. Figure 3 depicts an example of how the three schedulers behave. The system is simulated for a two year excerpt of the MD data and errors are drawn from $e(t) \sim \mathcal{N}(0.06, 0.05)$. The OPT algorithm determines a relatively stable and high used energy, the LT-ENO's used energy is lower but remains constant. Lastly, the DCA algorithm generates a used energy trace that has the most short term variability.

4.3 Robustness of the Minimal Used Energy

First, the robustness analysis is shown in detail for location MD and subsequently, the results for all locations are summarized. The confidence interval of the minimal used energy per day median for

(a) Prediction errors affect the minimal used energy per day.

Table 2: Robustness, R and performance, minimal used energy per day ($MinE$), analysis for all locations.

Schedulers & Robustness, Performance		Data set							
		AK	CA	MD	MI	ON	OR	TX1	TX2
OPT	$MinE$ [Wh]	1.50	3.54	3.65	2.88	2.61	2.47	3.98	3.43
	R	0.56	0.37	0.44	0.35	0.40	0.52	0.43	0.36
LT-ENO	$MinE$ [Wh]	1.37	2.08	2.10	1.76	1.70	2.17	3.12	2.14
	R	49.57	8.34	4.85	27.3	17.38	20.96	4.47	4.78
DCA	$MinE$ [Wh]	0.13	0.61	0.92	0.53	0.58	0.71	0.96	1.16
	R	1.07	0.64	0.33	0.42	0.38	0.34	0.38	0.27

location MD is depicted in Figure 4(a). The confidence interval spans a very small range, thus only the lower bound is shown and used for the analysis. The OPT algorithm has the highest performance, for zero systematic error the minimal used energy per day median is $MinE = 3.65$ Wh. However, its performance decreases with increasing prediction error median and (3) evaluates to $R = 0.44$. The performance of the LT-ENO algorithm is lower, $MinE = 2.1$ Wh, but it is hardly affected by prediction errors, the robustness metric is $R = 4.85$. The DCA algorithm has the lowest performance, $MinE = 0.92$ Wh, and is the least robust, $R = 0.33$.

Table 2 shows the results of this analysis for all the locations in Table 1. The OPT algorithm has the highest performance across all locations whereas the LT-ENO scheduler has the highest degree of robustness although at a lower performance level. Lastly, for this performance metric the DCA algorithm is strongly dominated by the other two schedulers.

Robustness Performance Trade-off The OPT scheduler is provably optimal with respect to the minimal used energy when predictions are accurate. It provides the highest performance as it aggressively exploits the predicted energy. However, when the environment differs from the predictions, the performance is severely impacted. Although the OPT scheduler is optimal for accurate predictions it does not portrair a high degree of robustness. Thus there is a trade-off between performance and robustness.

4.4 Robustness of the Total Utility

Figure 4(b) depicts the lower bound of the total utility median for location MD. Equivalent to the previous analysis, only the lower bound is shown and used for the robustness analysis. The total utility median achieved by the OPT and LT-ENO algorithm barely changes with varying prediction errors, $R = 59.6$ and $R = 19.35$ respectively. This is a result of the square root function only mildly penalizing high used energies. Conversely, the DCA algorithm displays a significantly lower degree of robustness, $R = 0.33$. Table 3 summarizes this analysis for all the locations in Table 1.

(b) Prediction errors affect the total utility.

Figure 4: Robustness performance analysis for the minimal used energy per day and total utility. The value at zero median error indicates the performance and the slope is related to the robustness.

Table 3: Robustness, R and performance, total utility ($TotalUtil$), analysis for all locations.

Schedulers &		Data set							
Robustness, Performance		AK	CA	MD	MI	ON	OR	TX1	TX2
OPT	$TotalUtil$	3083	3448	3613	3463	3623	3357	3635	3239
	R	47.1	53.3	59.6	51.4	52.8	53.6	65.0	58.4
LT-ENO	$TotalUtil$	1802	2523	2571	2364	2287	2263	2832	2443
	R	322.03	8.59	19.35	12.5	801.93	441.79	12.69	12.30
DCA	$TotalUtil$	1862	3157	3267	2777	2755	2754	3334	2992
	R	0.28	0.37	0.33	0.32	0.31	0.30	0.34	0.37

4.5 Various Prediction Error Distributions

We demonstrate that the error distribution does not significantly impact the results and conclusions of the robustness analysis. The analysis is performed for normal error distributions with different magnitudes of the random error, for left- and right-skewed normal distributions and for the heavy-tailed Cauchy distribution. Non-symmetric error distributions represent the case when it is more likely to largely overestimate as opposed to severely underestimate, or vice versa. Prediction errors from heavy-tailed distributions are more likely to be large. Table 4 summarizes the results for location MD. This evaluation indicates that the robustness analysis does not greatly depend on the prediction error distribution and thus supports the use of normally distributed errors.

5 IMPROVING THE DEGREE OF ROBUSTNESS

Various methods can be employed to improve a system's degree of robustness. Possible methods include scaling the energy prediction that the scheduler utilizes, scaling the scheduled energy or retaining part of the energy storage for critical situations.

5.1 Prediction Scaling Method

The robustness can be improved by providing the scheduler with a scaled energy prediction. Thus between the predictor and scheduler the prediction is scaled by $c \in (0, 1]$, as depicted in Figure 2. Scaling the prediction effectively changes the error the scheduler perceives. The prediction error is $E_{pred}(t) = E_{harv} \cdot (1 + e(t))$ whereas the error the scheduler perceives is $e'(t) = (c - 1) + c \cdot e(t)$. When c is chosen to be smaller than 1, the perceived error is scaled as well as moved to systematically underestimate the harvested energy. This typically increases the degree of robustness.

Table 4: Robustness R and performance for each of the metrics, minimal used energy per day ($MinE$) and total utility ($TotalUtil$).

Schedulers &		Error Distributions					
Performance		Normal	Normal	Normal	Left skewed	Right skewed	Cauchy
Robustness		$\sigma = 0.01$	$\sigma = 0.05$	$\sigma = 0.1$	$\sigma = 0.05$	$\sigma = 0.05$	$\sigma = 0.05$
OPT	$MinE$ [Wh]	3.83	3.65	3.43	3.73	3.53	3.63
	R	0.42	0.44	0.44	0.46	0.46	0.47
	$TotalUtil$	3614	3613	3612	3613	3613	3613
	R	57.0	59.6	63.3	64.8	55.4	56.9
LT-ENO	$MinE$ [Wh]	2.3	2.10	1.67	2.08	2.08	2.28
	R	9.54	4.85	0.83	4.66	2.50	5.78
	$TotalUtil$	2580	2571	2538	2566	2567	2581
	R	87.79	19.35	10.30	5.76	22.78	11.78
DCA	$MinE$ [Wh]	0.94	0.92	0.90	0.93	0.92	0.92
	R	0.35	0.33	0.28	0.34	0.31	0.29
	$TotalUtil$	3504	3267	3002	3347	3198	3444
	R	0.24	0.33	0.46	0.34	0.31	0.29

Figure 5: The performance and influence of prediction errors decreases with decreasing prediction scaling factor.

5.2 Case Study: Prediction Scaling

In Section 4, we concluded that although the OPT scheduler is optimal when predictions are accurate, it is not very robust. To improve the OPT scheduler's robustness, the prediction scaling method from Section 5.1 is employed. Implementation of the scaling method is resource efficient and only imposes a negligible overhead. Figure 5 shows the lower bound of the minimal used energy per day median for scaling factors $c \in [1, 0.95, 0.9]$.

Without scaling, $c = 1$, the performance is $MinE = 3.65$ Wh and the robustness metric $R = 0.44$. The robustness improves to $R = 0.69$ for a scaling factor of $c = 0.95$ and $R = 1.36$ for $c = 0.9$. It thus improves the degree of robustness. However, the performance changes to $MinE = 3.67$ Wh and decreases to 3.52 Wh for $c = 0.95$ and $c = 0.9$, respectively. The total utility is inadvertently also affected by the prediction scaling: its robustness improves and its performance decreases by less than 1 %. The above evaluation for multiple locations is summarized in Table 5.

Robustness Performance Trade-off In Section 4 a robustness performance trade-off was described. This trade-off can be explored with prediction scaling. Prediction scaling can improve the degree of robustness but the performance decreases when the environment indeed closely follows the predictions, as seen in Table 5. Thus the robustness performance trade-off can be tuned.

6 EXPERIMENTAL VALIDATION: INDOOR HARVESTING SCENARIO

The conclusions from Section 4 and 5.2 are validated in a real-world indoor harvesting scenario.

6.1 System and Data

Indoor harvesting data [22] spanning from August 2017 to June 2019 with a time interval of 1 hour for two employee offices is used. The first office (LOC 0) has little natural sunlight and thus the harvested energy is tied to office hours. The second office (LOC 1)

Table 5: Prediction scaling allows to explore the robustness performance trade-off.

Data sets, Performance & Robustness						
Scaling factors	CA		MD		MI	
	$MinE$ [Wh]	$TotalUtil$	$MinE$ [Wh]	$TotalUtil$	$MinE$ [Wh]	$TotalUtil$
	R	R	R	R	R	R
$c = 1$	3.54	3448	3.65	3613	2.88	3463
	0.37	53.3	0.44	59.6	0.35	51.4
$c = 0.95$	3.57	3447	3.67	3612	2.87	3462
	0.66	105.5	0.69	127.0	0.70	104.0
$c = 0.9$	3.42	3445	3.52	3610	2.76	3459
	1.19	71.1	1.36	68.3	1.37	67.2

Figure 6: Sample hourly harvested energy of two offices.

is exposed to significant levels of sunlight. Samples of the harvested energy can be seen in Figure 6.

The system consists of a solar panel with size $50 \text{ mm} \cdot 33 \text{ mm}$ and an energy storage with capacity $B = 3 \text{ J}$. The maximum power consumption is 0.1 mW and minimal duty cycle 0.01 . The three schedulers are combined with two energy predictors providing different levels of accuracy. The first predictor (CONST) predicts the harvested energy of each hour to be as much as was harvested in the previous hour. The second predictor is the exponentially-weighted moving average prediction model (EWMA) with $\alpha = 0.5$. The prediction accuracy is quantified by the mean absolute deviation percentage (MADP) as defined in [18]. The CONST predictor provides a higher prediction accuracy, the MADP is 38 % for LOC 0 and 56 % for LOC 1, than the EWMA algorithm, for which the MADP is 54 % and 73 % for LOC 0 and LOC 1, respectively.

6.2 Robustness of the Minimal Used Energy

Figure 7 depicts the minimal used energy per hour for each of the three schedulers when combined with either the CONST, high accuracy, predictor or the EWMA, low accuracy, predictor. The OPT algorithm has the highest minimal used energy per hour when the prediction accuracy is high because the scheduler is provable optimal when predictions are accurate and high accuracy predictions do not deviate too much from that scenario. However, its performance drops by at least an order of magnitude when the prediction accuracy is low. This is a result of the scheduler heavily relying and fully exploiting all predictions. Conversely, the LT-ENO scheduler has comparable performance for both predictors. This scheduler relies on sums of predictions and thus is more robust against prediction errors. Lastly, the performance of the DCA algorithm is lower than that of the other two schedulers and decreases with decreasing prediction accuracy. The robustness performance trade-off for the optimal OPT scheduler is visible in both locations. The OPT scheduler is so gravely impacted by the prediction errors, that the LT-ENO scheduler outperforms the OPT scheduler for the low accuracy predictors.

(a) Prediction scaling at LOC 0

(a) Robustness at LOC 0

(b) Robustness at LOC 1

Figure 7: In an indoor harvesting scenario the performance difference between high and low accuracy predictor aligns with the robustness analysis from Section 4.

6.3 Prediction Scaling

To alleviate the drastic performance drop of the OPT scheduler, we apply the prediction scaling method. The resulting minimal used energy per hour scheduled by the OPT algorithm for a range of prediction scaling factors is shown in Figure 8. By scaling the prediction with 0.3 for LOC 0 and 0.2 for LOC 1, the minimal used energy per hour determined by the OPT scheduler with the EWMA predictor can be improved by $24\times$ for LOC 0 and by $75\times$ for LOC 1. However, prediction scaling reduces the performance for the high accuracy case, where for LOC 0, the minimal used energy decreases by 33 % and for LOC 1 by 12 %. Thus the trade-off for the OPT algorithm is illustrated and tuned with prediction scaling.

7 CONCLUSION

Energy harvesting systems with minimum performance guarantees are challenging to design due to the non-deterministic nature of their energy source. Furthermore, the interdependence between energy predictors, schedulers and system performance complicates their optimization such that optimality is only achieved in a limited range of scenarios. The presented robustness analysis identifies these scenarios by studying the effects of prediction accuracy on system performance. We studied different combinations of existing energy predictors and schedulers that are either optimal or robust. This demonstrates the low degree of robustness of an optimal scheduler and thus a robustness performance trade-off. Our results show prediction scaling can improve worst-case system performance by a factor of 75. Robustness can be used to evaluate systems and identify cases where additional measures such as the prediction scaling method need to be taken to achieve the desired system behavior in a wide range of environments. As future work, we will investigate whether the scaling factor can be adjusted online.

(b) Prediction scaling at LOC 1

Figure 8: Prediction Scaling increases the minimal used energy per hour by $24\times$ and $75\times$ for LOC 0 and LOC 1, respectively, when a low accuracy predictor is used. The performance penalty when using the high accuracy predictor is only 33 % for LOC 0 and 12 % for LOC 1.

REFERENCES

- [1] Aman Kansal, Jason Hsu, Sadaf Zahedi, and Mani B. Srivastava. Power Management in Energy Harvesting Sensor Networks. *ACM Trans. Embed. Comput. Syst.*, 6(4), 2007.
- [2] Selahattin Kosunalp. A New Energy Prediction Algorithm for Energy-Harvesting Wireless Sensor Networks With Q-Learning. *IEEE Access*, 4:5755–5763, 2016.
- [3] Hengzhao Yang and Ying Zhang. A task scheduling algorithm based on supercapacitor charge redistribution and energy harvesting for wireless sensor nodes. *Journal of Energy Storage*, 6:186–194, 2016.
- [4] Antonio Caruso, Stefano Chessa, Soledad Escobar, Xavier Del Toro, and Juan Carlos López. A dynamic programming algorithm for high-level task scheduling in energy harvesting iot. *IEEE Internet of Things Journal*, 5(3):2234–2248, 2018.
- [5] Abdulmajid Murad, Frank Alexander Kraemer, Kerstin Bach, and Gavin Taylor. Autonomous management of energy-harvesting iot nodes using deep reinforcement learning. In *2019 IEEE 13th International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, pages 43–51. IEEE, 2019.
- [6] Rehan Ahmed, Bernhard Buchli, Stefan Draskovic, Lukas Sigrist, Pratyush Kumar, and Lothar Thiele. Optimal power management with guaranteed minimum energy utilization for solar energy harvesting systems. *ACM Trans. Embed. Comput. Syst.*, 18(4), 2019.
- [7] P. Anagnostou and A. Gomez and P.A. Hager and H. Fatemi and J. Pineda de Gyvez and L. Thiele and L. Benini. Energy and power awareness in hardware schedulers for energy harvesting IoT SoCs. *Integration*, 67:33 – 43, 2019.
- [8] Bernhard Buchli, Pratyush Kumar, and Lothar Thiele. Optimal power management with guaranteed minimum energy utilization for solar energy harvesting systems. In *2015 International Conference on Distributed Computing in Sensor Systems*, pages 147–158. IEEE, 2015.
- [9] Bernhard Buchli, Felix Sutton, Jan Beutel, and Lothar Thiele. Dynamic power management for long-term energy neutral operation of solar energy harvesting systems. In *Proceedings of the 12th ACM conference on embedded network sensor systems*, pages 31–45. ACM, 2014.
- [10] Roy Chaoming Hsu and Tzu Hao Lin. A Fuzzy Q-learning Based Power Management for Energy Harvest Wireless Sensor Node. In *2018 International Conference on High Performance Computing Simulation (HPCS)*, pages 957–961. IEEE, 2018.
- [11] Yue Xu, Hyung Gyu Lee, Yujuan Tan, Yu Wu, Xianzhang Chen, Liang Liang, Lei Qiao, and Duo Liu. Tumbler: Energy efficient task scheduling for dual-channel solar-powered sensor nodes. In *2019 56th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6, 2019.
- [12] Bruno Srbinovski, Michele Magno, Fiona Edwards-Murphy, Vikram Pakrashi, and Emanuel Popovici. An energy aware adaptive sampling algorithm for energy harvesting WSN with energy hungry sensors. *Sensors*, 16(4), 2016.
- [13] Hassaan Khaliq Qureshi, Umber Saleem, Muhammad Saleem, Andreas Pitsillides, Marios Lestas, et al. Harvested energy prediction schemes for wireless sensor networks: Performance evaluation and enhancements. *Wireless Communications and Mobile Computing*, 2017.
- [14] Murad Al-Omary, Khaoula Hassini, Ahmed Fakhfakh, and Olfa Kanoun. Prediction of Energy in Solar Powered Wireless Sensors Using Artificial Neural Network. In *2019 16th International Multi-Conference on Systems, Signals Devices (SSD)*, pages 288–293. IEEE, 2019.
- [15] Zhaohui Zhang, Sanyang Liu, Yanqi Yang, and Yiguang Bai. A link-adding strategy for improving robustness and traffic capacity in large-scale wireless sensor networks. *Cluster Computing*, 22(3):7687–7694, 2019.
- [16] Alma A.M. Rahat, Richard M. Everson, and Jonathan E. Fieldsend. Evolutionary multi-path routing for network lifetime and robustness in wireless sensor networks. *Ad Hoc Networks*, 52:130–145, 2016.
- [17] Orhan Dagdeviren and Vahid Khalilpour Akram. The effect of random node distribution and transmission ranges on connectivity robustness in wireless sensor networks. In *2019 International Symposium on Networks, Computers and Communications (ISNCC)*, pages 1–5. IEEE, 2019.
- [18] Eva Lucas Segarra, Hu Du, Germán Ramos Ruiz, and Carlos Fernández Banderá. Methodology for the quantification of the impact of weather forecasts in predictive simulation models. *Energies*, 12(7), 2019.
- [19] William R. Thompson. On Confidence Ranges for the Median and Other Expectation Distributions for Populations of Unknown Distribution Form. *The Annals of Mathematical Statistics*, 7(3):122–128, 1936.
- [20] http://rredc.nrel.gov/solar/old_data/nsrdb/.
- [21] Bernhard Buchli, Felix Sutton, Jan Beutel, and Lothar Thiele. Towards enabling uninterrupted long-term operation of solar energy harvesting embedded systems. In *European Conference on Wireless Sensor Networks*, pages 66–83. Springer, 2014.
- [22] Lukas Sigrist, Andres Gomez, and Lothar Thiele. Dataset: Tracing indoor solar harvesting. In *Proceedings of the 2nd Workshop on Data Acquisition To Analysis*, page 47–50. ACM, 2019.