

# Dynamic Power Management for Long-Term Energy Neutral Operation of Solar Energy Harvesting Systems

Bernhard Buchli, Felix Sutton, Jan Beutel, Lothar Thiele  
Computer Engineering and Networks Laboratory  
Swiss Federal Institute of Technology (ETH) Zurich  
Zurich, Switzerland  
{bbuchli, fsutton, beutel, thiele}@tik.ee.ethz.ch

## Abstract

In this work we consider a real-world environmental monitoring scenario that requires uninterrupted system operation over time periods on the order of multiple years. To achieve this goal, we propose a novel approach to dynamically adjust the system's performance level such that energy neutral operation, and thus long-term uninterrupted operation can be achieved. We first consider the annual dynamics of the energy source to design an appropriate power subsystem (*i.e.*, solar panel size and energy store capacity), and then dynamically compute the long-term sustainable performance level at runtime. We show through trace-driven simulations using eleven years of real-world data that our approach outperforms existing predictive, *e.g.*, EWMA, WCMA, and reactive, *e.g.*, ENO-MAX, approaches in terms of average performance level by up to 177%, while reducing duty-cycle variance by up to three orders of magnitude. We further demonstrate the benefits of the dynamic power management scheme using a wireless sensor system deployed for environmental monitoring in a remote, high-alpine environment as a case study. A performance evaluation over two years reveals that the dynamic power management scheme achieves a two-fold improvement in system utility when compared to only applying appropriate capacity planning.

## Categories and Subject Descriptors

C.4 [Performance of Systems]: Design studies, modeling techniques

## General Terms

Algorithms, design, experimentation

### Keywords

Energy neutral operation, solar energy harvesting, wireless sensor networks

## 1 Introduction

The performance level achievable by an embedded system is ultimately limited by the energy available to operate the device. Due to the predominantly remote deployment locations of Wireless Sensor Networks (WSNs) and the lack of dependable power sources, the motes comprising these networks generally rely on batteries for delivering the energy to fulfill their intended task. However, due to the finite capacity of the energy storage element, *i.e.*, battery, the motes are highly energy constrained and suffer from a severely constrained lifetime. To improve the system's achievable performance level and extend the lifetime, ambient energy harvesting, particularly in the form of solar energy harvesting, has been established as a feasible alternative to purely battery powered devices in outdoor WSN applications [20].

Using a real-world application scenario [3], which requires high system availability, and relies on sensing technology characterized by high energy demands, we investigate if the two conflicting goals, *i.e.*, high system performance and lifetime on the order of multiple years, can be simultaneously achieved with solar energy harvesting systems. A broad range of application scenarios, *e.g.*, [9, 11, 22, 24], benefit from a minimum supported performance level that can be sustained over time periods on the order of multiple years. A system enhanced with energy harvesting capabilities can – in theory – operate indefinitely as the energy store can be replenished periodically. Experience has shown, however, that enhancing a battery operated device with energy harvesting capabilities will by itself neither provide a lower bound on the expected sustainable performance level, nor guarantee uninterrupted long-term operation [21]. The reason for this is the dependence on an uncontrollable energy source [15], *i.e.*, the sun, which exhibits high short-term fluctuations due to meteorological conditions that are hard to model [8] and difficult to predict [13].

Contemporary power management techniques deal with the highly variable energy harvesting opportunities by dynamically adapting the system's performance level at runtime such that Energy Neutral Operation (ENO) [15] may be achieved. Informally, a system is said to operate in an energy neutral mode if the energy consumed over a given time period  $\delta$  is less than or equal to the energy harvested during the same time period. Due to practical limitations, ENO is generally interpreted such that the battery fill-level  $B_{fill}$  at

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [Permissions@acm.org](mailto:Permissions@acm.org).

SenSys'14, November 3–5, 2014, Memphis, TN, USA.  
Copyright 2014 ACM 978-1-4503-3143-2/14/11 ...\$15.00  
<http://dx.doi.org/10.1145/2668332.2668333>

the end of period  $\delta$  must be greater than or equal to that at the beginning, *i.e.*,  $B_{fill}(t + \delta) \geq B_{fill}(t)$  [15, 18, 23].

Given ENO as the fundamental bound of energy harvesting systems, numerous methods that attempt to achieve this objective have been proposed, *e.g.*, [15–18, 23]. These can be classified as (i) predictive, and (ii) reactive approaches. Predictive approaches, *e.g.*, [15, 18], attempt to satisfy ENO by predicting the harvestable energy during a future time slot, and adapt the performance level accordingly. However, predicting future meteorological conditions is highly complex and may be computationally prohibitive [8]. Therefore, acceptable prediction accuracy with the limited computational resources available on contemporary motes has so far only been possible for short prediction windows, *i.e.*,  $\delta$  on the order of minutes to hours.

Reactive approaches, on the other hand, attempt to satisfy energy neutrality by scheduling the performance level in response to changes in the source. This can be done by measuring the energy generation directly, or, as is commonly done, through monitoring the battery fill-level [23], or super-capacitor voltage [16]. The performance of a storage-reactive approach strongly relies on the accuracy of the battery State-of-Charge indication.

Current implementations of the above two classes adapt the system duty-cycle in response to, or expectation of, short-term variations of the energy source, and thus tend to suffer from high duty-cycle variance. Duty-cycle variance is an important consideration, *e.g.*, for surveillance applications, where the system should be available with equal probability at any given point in time [12].

In this work we turn our attention to enabling long-term energy neutral operation for solar energy harvesting systems. Rather than predicting or reacting to the source’s *short-term* variations, we argue that the source’s *long-term* dynamics must be considered both for dimensioning the power subsystem and devising the dynamic energy management scheme. We leverage the approach discussed in [7] to provision the power subsystem, *i.e.*, battery and solar panel, such that short-term fluctuations can be absorbed. We further devise a long-term energy-predictive dynamic power management technique that can compute the long-term sustainable performance level at runtime.

The contributions of this work are as follows. First, we present an end-to-end solution for enabling Long-Term Energy Neutral Operation (LT-ENO) for solar energy harvesting systems. Our approach encompasses (i) a power subsystem capacity planning algorithm based on an astronomical solar radiation model, and (ii) a dynamic energy management scheme, which is based on the same astronomical model, and that can enable uninterrupted operation with very low duty-cycle variance. Second, through simulation with eleven years of data at three different geographical locations, we show that our algorithm outperforms the State-of-the-Art in energy-predictive [15, 18], and battery-reactive [23] performance scaling approaches in terms of average sustainable performance level by up to 177%, energy efficiency by up to 184%, and duty-cycle stability by up to three orders of magnitude, while incurring zero downtime, *i.e.*, system availability of 100%. Finally, we exemplify the benefits of

our approach using an X-SENSE environmental monitoring system [3] deployed over two years in a high-alpine environment, and demonstrate that significant improvements in system utility can be achieved without risking downtime due to power outages.

The rest of this paper is structured as follows. *Sec. 2* reviews the State-of-the-Art in power subsystem capacity planning, energy prediction schemes, and harvesting aware dynamic performance scaling techniques. *Sec. 3* reviews the power subsystem capacity planning approach. The dynamic power management technique for LT-ENO is discussed in detail in *Sec. 4*. In *Sec. 5* the proposed technique is evaluated through simulation with eleven years of data for three different locations, while *Sec. 6* provides performance results obtained with a system deployed over a period of two years. *Sec. 7* concludes this work with a summary of key findings.

## 2 Related Work

**Capacity Planning.** The importance of proper capacity planning for solar energy harvesting systems has been introduced in [15], and a systematic technique was proposed. The approach relies on the availability of a representative energy generation profile and the known system consumption to compute the battery capacity. The limitations of this approach are two-fold. First, the input trace, *i.e.*, energy profile, must be representative of the conditions at the intended deployment site, and cover at least one full annual solar cycle to yield a suitable battery capacity. Second, the panel size is not considered a design parameter, thus preventing the designer from optimizing the power subsystem with respect to cost, physical form-factor, *etc.*

A recently proposed approach [7] mitigates the aforementioned shortcomings. The authors propose a capacity planning algorithm that relies on an astronomical model to approximate the energy profile at the intended deployment site. With this approach, the designer can vary all important design parameters to obtain the specifications of a suitable power subsystem for a given application without the need for extensive trace data. Through simulation with ten years of trace data, it was shown that the power subsystem obtained enables uninterrupted operation if the actual total energy generation is at least 80% of the modeled expectations. In this work we develop a dynamic power management scheme that builds upon this capacity planning technique.

**Energy Management.** In the seminal work on energy harvesting theory [15], the first dynamic duty-cycling scheme for solar energy harvesting systems was proposed within a theoretical framework that defines Energy Neutral Operation (ENO) as the fundamental limit of energy harvesting systems. ENO is achieved if the system never consumes more energy than what it can harvest over a given time period  $\delta$ , *i.e.*, the battery fill-level  $B_{fill}(t + \delta)$  is greater than or equal to  $B_{fill}(t)$ . With their approach, a day is discretized into slots of equal duration  $\delta$ , and the expected energy input for each slot is learned with an Exponentially Weighted Moving Average (EWMA) filter. Each slot’s respective duty-cycle is then computed by considering the mismatch between expected and actual energy input. However, due to limited correlation between past and future weather conditions, this approach

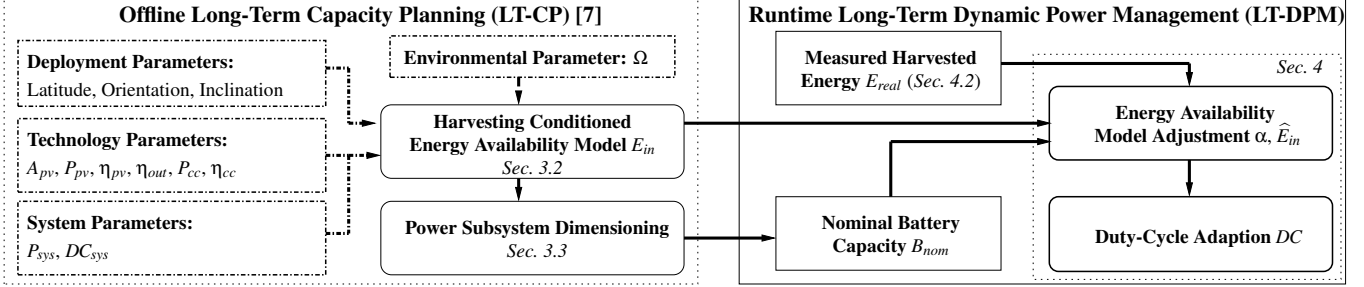


Figure 1: Process flow for long-term solar energy harvesting capacity planning and dynamic power management. Dashed boxes and arrows represent user inputs. The offline capacity planning algorithm computes the achievable duty-cycle and required battery capacity for the given input parameter set, and the dynamic power management algorithm adjusts the system performance level at runtime according to the observed conditions.

achieves acceptable prediction accuracy only for prediction windows on the order of hours.

Weather Conditioned Moving Average (WCMA), proposed in [18], improves upon EWMA’s prediction accuracy. The authors not only consider the harvested energy in the same time slot during previous days, but also incorporate current weather conditions to obtain the expected energy input in the current slot. While achieving an almost three-fold improvement in prediction accuracy over EWMA, it is not clear if and how this improvement translates into increased system performance and/or energy neutrality. This approach is also constrained by short prediction windows.

More recently, the use of professional weather forecast services have been considered to predict the disposable energy [19]. The authors formulate a model to translate weather forecasts into solar or wind energy harvesting predictions. While it is unclear what baseline is used, the authors conclude that their energy predictions are more accurate than those based on past local observations.

In [16] and [23], model-free approaches to dynamic performance scaling are presented. In [23], a technique from adaptive control theory, *i.e.*, Linear-Quadratic Tracking, is used to dynamically adapt the system’s duty-cycle based on the battery State-of-Charge and so ensure ENO. For the datasets evaluated, the authors report between 6 and 32% improvement in mean duty-cycle, and between 6 and 69% reduction in duty-cycle variance when compared to EWMA. Similarly, in [16] a Proportional-Integral-Derivative (PID) controller monitors the energy storage element, and the duty-cycle is adapted such that an expected voltage level of the storage element (a super-capacitor in this case) is maintained. While presenting low-complexity solutions, both of these approaches suffer from high duty-cycle variability, and rely on a well performing battery State-of-Charge approximation algorithm. The PID approach additionally requires parameter tuning, for which solutions exist in the literature.

### 3 Capacity Planning for Long-Term Energy Neutral Operation

Rather than modeling the energy source’s highly variable short-term dynamics and adjust the performance level accordingly, we propose a long-term energy neutral power management scheme for solar energy harvesting systems. Our approach, illustrated in *Figure 1*, first invokes a design-time power subsystem capacity planning algorithm to deter-

mine the required battery capacity given a set of input parameters that characterize the system and its environment. The intricate trade-offs between battery capacity, and the system and environmental parameters are discussed in [7]. This algorithm uses an astronomical model to estimate the long-term energy availability based on the annual solar cycle. Then, at runtime, the proposed algorithm dynamically computes the performance level, *i.e.*, duty-cycle, based on an adjusted energy availability model such that long-term energy neutrality can be sustained. The energy model and the capacity planning approach follow [7] and are briefly reviewed in this section. The novel dynamic power management scheme is discussed in detail in *Sec. 4*, and evaluated in *Sec. 5* and *6*.

#### 3.1 System Architecture, Load Model, and System Utility

In this work we assume a harvest-store-use architecture, as described in [20], in which the energy to operate the system is always supplied by the battery. We further assume that the power  $P_{sys}$  dissipated from the battery includes all consumers present in the system, *e.g.*, power conditioning and other supervisory circuitry. Further considering that contemporary embedded systems can operate in sleep modes with ultra-low power dissipation, we ignore its contribution and define the total daily energy  $E_{out}(d)$  necessary to sustain a required performance level  $DC_{sys}(d)$  on calendar day  $d$  as given in (1), where  $\gamma = 24 \text{ hours}$ . Note that we ignore battery leakage here, but it can be integrated into the load model.

$$E_{out}(d) = \gamma \cdot DC_{sys}(d) \cdot P_{sys}, \quad \forall d \in \mathbb{Z}^+ \quad (1)$$

For now, we assume a one-to-one relationship between performance level  $DC_{sys}(d)$  and utility of the system  $U$ , *i.e.*,  $U(DC_{sys}(d)) = DC_{sys}(d)$  [9]. We revisit this topic in *Sec. 6*, where we refine the definition of system utility in the context of a real system. Note that we are not concerned with how the energy is scheduled and consumed over the course of the day, but rather provide information about disposable energy to an application specific task scheduler. Details on local scheduling of the available energy, and network-wide balancing of the energy budget by changing the communication and/or sensing patterns are beyond the scope of this paper, as they are highly application specific. For example, a scheduler’s primary focus may be planning the available energy such that a minimum level of operation may be sustained. Any excess energy may then be used to improve sensing,




Figure 2: Solar energy profile for a particular geographical location and energy harvesting setup. Surplus energy generated by the panel is indicated with the hatched area; the energy deficit is shown by the cross-hatched area.

processing or communication.

### 3.2 Harvesting Conditioned Energy Availability Model

A crucial step in capacity planning consists of estimating the theoretically harvestable energy at a specific point in space and time. Figure 2 illustrates the amount of solar energy harvested at a particular geographical location and given harvesting configuration. The figure shows the total daily energy input  $E_{real}(d)$  at the end of each calendar day  $d$ , and illustrates the high short-term (day-to-day) variability and long-term periodicity (year-to-year) of the source. Also shown is the modeled total expected harvestable energy  $E_{in}(d)$  on calendar day  $d$  such that true energy conditions are closely approximated, *i.e.*, (2) holds where  $N$  is the number of days.

$$\sum_{d=1}^N E_{in}(d) \cong \sum_{d=1}^N E_{real}(d), \quad N \gg 1 \quad (2)$$

The method to compute  $E_{in}(d)$  is based on a simplified astronomical model to estimate the theoretical solar radiation  $E_{astro}(t, d, L, \theta_p, \phi_p, \Omega)$ . It is parameterized by the time  $t$  in hours of calendar day  $d$ , the intended deployment site's latitude  $L$ , and the panel's orientation and inclination angles  $\phi_p$  and  $\theta_p$ , respectively. Finally, the environmental parameter  $\Omega$  represents the expected average meteorological conditions. This is the only unknown input parameter, and can be approximated as described in [7]. Although not absolutely necessary, the availability of solar maps or solar energy traces can improve the approximation of the parameter  $\Omega$ .

Since we are concerned with electrical, as opposed to solar energy, the output of  $E_{astro}(\cdot)$  must be conditioned by the technology parameters in Figure 1. These specify the panel's surface area  $A_{pv}$ , conversion efficiency  $\eta_{pv}$ , and self-consumption and efficiency factors for supervisory and power conditioning circuitry, *e.g.*, battery charge controller efficiency  $\eta_{cc}$ , and consumption  $P_{cc}$ . The maximum rated power output of the panel  $P_{pv}$  is used to evaluate the maximum energy  $E_{pv}$  generated during one hour. Then, with the above parameters specified, the total electrical energy that can be harvested on calendar day  $d$  is approximated with (3).

$$E_{in}(d) = A_{pv} \eta_{cc} \eta_{pv} \sum_{t=1}^{24} \min(E_{pv}, E_{astro}(t, d, \dots)) \quad (3)$$

While the astronomical energy model  $E_{astro}(\cdot)$  may yield any resolution  $t$ , for the purpose of long-term energy neutral operation discussed in this work, daily sums are sufficient.

### 3.3 Power Subsystem Dimensioning

In this section we review the process of computing the power subsystem capacity using the energy availability model such that energy neutral operation over the source's seasonal cycle, *i.e.*, one year, can be achieved. At this point we assume a perfect battery, *i.e.*, no inefficiencies. For a discussion including various battery inefficiencies, the reader is referred to [7].

For the purpose of power subsystem capacity planning we assume a constant daily energy demand  $E_{out}(d)$  that must be met. Note that we explicitly keep the dependence on calendar day  $d$ , since the energy consumption at runtime varies with the dynamically chosen daily duty-cycle (see Sec. 4). Referring to Figure 2, we observe that the intersections between the energy consumption  $E_{out}(d)$  and approximated energy input  $E_{in}(d)$  partition the annual solar cycle into time regions of energy surplus, *i.e.*,  $E_{in}(d) > E_{out}(d) \forall d \in [d_0, d_1]$ , and energy deficit, *i.e.*,  $E_{in}(d) < E_{out}(d) \forall d \in [d_1, d_2]$ .

According to the model assumptions, the minimum battery capacity  $B$  required to support the system during periods of energy deficit is indicated with the cross-hatched area in Figure 2, and formally stated in (4). The first term on the left-hand side defines the amount of energy that is necessary to support the system operation, while the second term represents the expected energy input. The difference is then the minimum required battery capacity.

$$\sum_{d_1}^{d_2} (E_{out}(d) - E_{in}(d)) \leq B \quad (4)$$

In order to achieve uninterrupted operation over multiple years, it is not sufficient to only provision the battery for the period of deficit. The panel must be able to generate enough energy to recharge the battery in addition to the energy required to sustain operation during periods of energy surplus, *i.e.*,  $d \in [d_0, d_1]$ . The constraint on energy generation by the panel is given in (5).

$$\sum_{d_0}^{d_1} (E_{in}(d) - E_{out}(d)) \geq B \quad (5)$$

The required battery capacity  $B$  can then be obtained by varying the performance level (*i.e.*,  $DC_{sys}(d)$ ) and/or the panel area  $A_{pv}$  and finding the intersections  $d_0$ ,  $d_1$ , and  $d_2$  between  $E_{in}(d)$  and  $E_{out}(d)$  such that (4) and (5) hold.

## 4 Dynamic Power Management for Long-Term Energy Neutral Operation

In the previous section we described the design-time energy availability model and power subsystem capacity planning based on the long-term characteristics of the energy source. Assuming that the design-time model reflects the conditions at the deployment location to within some bounds, the system will be able to run at the performance level for which the power subsystem was designed [7]. However, in practice significant deviations from the model must be expected. Such deviations may be caused by transient phenomena, *e.g.*, snow cover and foliage, or persistent occlusions due to trees and buildings. In this section we propose a dynamic power management scheme that can adapt

to deviations from the modeled assumptions by dynamically scaling the system performance level, and by doing so enable Long-Term Energy Neutral Operation (LT-ENO).

#### 4.1 Dynamic Performance Scaling

As discussed in *Sec. 3.3*, in order to achieve long-term energy neutrality, the two constraints from (4) and (5) must be satisfied. The constraint in (4) states that the battery must be able to supply the difference in energy consumption and generation during periods of energy deficit, *i.e.*,  $d \in [d_1, d_2]$  (as shown in *Figure 2*). The second constraint states that, in order to ensure that the battery can be fully recharged during periods of energy surplus ( $d \in [d_0, d_1]$ ), the panel must generate energy in excess of what is required to sustain short-term operation. To satisfy these two constraints, we leverage the offline energy model to determine the sustainable system performance level.

To exemplify our approach we consider a concrete example as illustrated in *Figure 3*. Without loss of generality, we assume that the design-time model  $E_{in}(d)$ , which was used to obtain the battery capacity  $B$  given panel size  $A_{pv}$ , overestimates the actual energy conditions  $E_{real}(d)$ . For simplicity we ignore battery inefficiencies in this discussion, but note that *Algorithm 1* and the evaluation in *Sec. 5* account for these effects. In the following we consider the end of day  $d$  and wish to compute the duty-cycle for the entire day  $d+1$  such that long-term energy neutrality may be achieved.

To react to deviations from the modeled energy expectation, we first need to adjust the design-time energy model  $E_{in}(d)$  at runtime according to observed conditions. For this purpose, we define the model adjustment factor  $\alpha$  in (6) to scale  $E_{in}(d)$ , *i.e.*,  $\hat{E}_{in}(d) = \alpha E_{in}(d)$ ,  $\forall d$ . The adjustment factor depends on the history window size  $W$  in days, which is used to tune the duty-cycle stability. The choice of  $W$  has a direct impact on the system's responsiveness to variations in the energy profile, and therefore imposes a system trade-off between duty-cycle stability and achievable performance level. The effects of the choice of the history window size  $W$  are discussed in *Sec. 5.3.3*.

$$\alpha = \frac{\sum_{d-W}^d E_{real}(d)}{\sum_{d-W}^d E_{in}(d)}, \quad 0 < W \leq d \quad (6)$$

Then, referring to *Figure 3*, it is evident that, given  $B$  and the adjusted energy model  $\hat{E}_{in}(d)$ , the modeled consumption  $E_{out}(d)$  may not be sustained. For example, a battery capacity dimensioned for  $d_1^*$ , and  $d_2^*$  instead of  $d_1$  and  $d_2$  would be necessary to support  $E_{out}(d)$  in *Figure 3*. Therefore, to fully, but safely leverage the available battery capacity given  $\hat{E}_{in}(d) < E_{in}(d) \forall d$ , we need to find the energy consumption  $\hat{E}_{out}(d) = DC(d) \cdot P_{sys} \cdot \gamma$ , where  $\gamma = 24$  hours such that the battery and panel constraints in (7) hold.

$$\sum_{d_1'}^{d_2'} \left( \hat{E}_{out}(d) - \hat{E}_{in}(d) \right) \leq B \leq \sum_{d_0'}^{d_1'} \left( \hat{E}_{in}(d) - \hat{E}_{out}(d) \right) \quad (7)$$

$$\begin{aligned} \alpha &= \frac{\sum_{d-W}^d E_{real}(d)}{\sum_{d-W}^d E_{in}(d)}; \\ \hat{E}_{in} &= \alpha \cdot \eta_{cc} \cdot E_{in}; \\ d_1' &= d_1; d_2' = d_2; d_0' = d_0; \\ surplus &= \sum_{d_0'}^{d_1'} \hat{E}_{in}(d) - \left( \hat{E}_{in}(d_1') \cdot (d_1' - d_0' + 1) \right); \\ deficit &= \left( \hat{E}_{in}(d_1') \cdot (d_2' - d_1' + 1) \right) - \sum_{d_1'}^{d_2'} \hat{E}_{in}(d); \\ \text{while } ((deficit \leq surplus) \ \&\& \ (surplus \leq B_{nom})) \ \text{do} \\ & \quad \text{if } \alpha < 1 \ \text{then} \\ & \quad \quad | \ d_0' = d_0' + 1; d_1' = d_1' - 1; d_2' = d_2' + 1; \\ & \quad \quad \text{else} \\ & \quad \quad | \ d_0' = d_0' - 1; d_1' = d_1' + 1; d_2' = d_2' - 1; \\ & \quad \quad \text{end} \\ & \quad \quad surplus, deficit = \text{calculate as above}; \\ & \quad \quad \text{if } (surplus < deficit) \ \text{then} \\ & \quad \quad \quad | \ d_0', d_1', d_2' = \text{previous } d_0', d_1', d_2'; \\ & \quad \quad \quad | \ surplus = \text{calculate as above}; \\ & \quad \quad \quad | \ break; \\ & \quad \quad \text{end} \\ \text{end} \\ DC(d+1) &= \frac{\min(B, surplus) + \sum_{d_1'}^{d_2'} \hat{E}_{in}(d)}{P_{sys} \cdot 24 \cdot (d_2' - d_1' + 1)}; \end{aligned}$$

**Algorithm 1:** Computation of the duty-cycle for day  $d+1$  performed at end of day  $d$ . In this example we use a daily resolution, which can be adapted to other time steps. Note that battery charge ( $\eta_{cc}$ ) and discharge efficiencies ( $\eta_{out}$ ) are incorporated, and nominal capacity  $B_{nom} = B/\eta_{out}$ .

The limits of summation in (7) are unknown and depend on  $\hat{E}_{out}(d)$ , the quantity we wish to find. However, since the modeled limits are known, or can be computed at runtime,  $d_0', d_1'$ , and  $d_2'$  can be found iteratively in discrete time steps, *e.g.*, days, starting with intervals  $D_s = [d_0, d_1]$ ,  $D_d = [d_1, d_2]$ , which represent the surplus and deficit regions respectively, and adjusting them according to *Algorithm 1* until (8) evaluates true.  $D_s^0$  and  $D_d^0$  in (8) denote the first elements in the intervals  $D_s$  and  $D_d$  respectively.

$$|D_d| \hat{E}_{in}(D_d^0) - \sum_{D_d} \hat{E}_{in}(d) \leq \sum_{D_s} \hat{E}_{in}(d) - |D_s| \hat{E}_{in}(D_s^0) \leq B \quad (8)$$

The relation in (8) is obtained from (7) by noting that, under our model assumptions, the energy generation at the start of the deficit period is equal to the consumption on that day (see *Figure 3*). Since we assume a constant energy consumption (*i.e.*, a stable duty-cycle) is desirable, we substitute  $\hat{E}_{in}(D_d^0)$  and  $\hat{E}_{in}(D_s^0)$  respectively, for  $\hat{E}_{out}(d)$  in (7), and replace the summations by multiplications.

Note that the maximum battery size that can be supported given the observed energy conditions is limited by the energy that can be harvested during the surplus period, *i.e.*,  $B \leq E_{surplus} = \left( \sum_{D_s} \hat{E}_{in}(d) \right) - \hat{E}_{in}(d_1') \cdot |D_s|$ . Taking this limitation into consideration, we can use (9) to compute the sustainable performance level for day  $d+1$  at the end of day  $d$ . Note,  $\gamma = 24$  hours.

$$DC(d+1) = \frac{\min(B, E_{surplus}) + \sum_{D_d} \hat{E}_{in}(d)}{P_{sys} \cdot \gamma \cdot |D_d|} \quad (9)$$

In summary, with the adjusted energy model we can approximate the expected energy input over the annual solar




Figure 3: Example where energy availability model  $E_{in}(d)$  overestimates the actual energy conditions  $E_{real}(d)$ . The performance level for  $d_{now} + 1$  is computed at the end of day  $d_{now}$ , using  $\hat{E}_{in}(d)$  with  $\alpha$  computed over the past  $W$  days.

cycle according to recent conditions. This information is used to continually adjust the long-term sustainable performance supported by the power subsystem. In other words, to ensure that the battery can be replenished during periods of surplus, and adequately used during periods of deficit, the performance level is computed by considering a full annual solar cycle. Note, in the above discussion we assumed that the design-time model overestimates true conditions. However, the approach, as shown in *Algorithm 1* is equally applicable to model underestimation.

## 4.2 Practical Considerations and Limitations

In this section we discuss implementation specific considerations and limitations of the proposed approach.

### 4.2.1 Measurement Support

The proposed dynamic power management scheme requires that the system can measure or approximate the total daily harvested energy. This can be accomplished by measuring the power output by the panel, or inferring the harvested energy through battery State-of-Charge information. The former is the preferred choice, but incurs additional overhead in terms of measurement circuitry and continual processing. The advantage of a State-of-Charge approach, e.g., [6] is that it may not require special purpose hardware, and needs to be performed only once a day. In *Sec. 6*, we show that it is indeed possible to use the technique discussed herein, even when the energy generated by the solar panel can not be measured directly.

### 4.2.2 Global Time Knowledge

Clearly, the proposed technique requires knowledge of global time in order to determine the current calendar day  $d$ . Considering that our approach achieves long-term energy neutrality, and may therefore operate without interruption (see *Sec. 5.2* and *6.2*), this is not considered a limitation.

### 4.2.3 Battery Inefficiencies

Batteries are non-ideal storage elements, which suffer from a variety of inefficiencies that are dependent on the specific battery chemistry and load behavior [2]. In our model, *charging* and *discharging* inefficiencies are incorporated through  $\eta_{cc}$  and  $\eta_{out}$  respectively specified by the system designer (see *Figure 1*). *Leakage power* is ignored in this discussion. Considering the periodically recurring recharging opportunities, accounting for leakage is not as crucial as it is for purely battery operated devices. *Temperature* may

impact the battery's apparent capacity [6]. Thus, for deployments that are exposed to low temperatures over extended periods of time, it may be necessary to account for the temporarily reduced battery capacity imposed by temperature effects. Finally, *battery aging* is not likely to be a problem, since batteries are generally rated for a few hundred deep discharge cycles [2]. With our approach, the battery experiences only one deep discharge cycle per year, and is therefore expected to outlast the lifetime of other system components, e.g., electronics, mechanical parts, etc. Note that the solar panel may also experience degradation. However, it has been shown that this tends to be aesthetic in nature, and does not significantly affect the panel's efficiency [10].

### 4.2.4 Worst Case Energy Conditions

The proposed energy neutral dynamic power management approach relies exclusively on the solar energy profile. Under normal circumstances, this is not a problem, as the duty-cycle is adapted according to the long-term dynamics of the source. However, in the case of a prolonged lack of harvesting opportunities, e.g., due to snow cover, the battery should be dimensioned such that this period can be bridged. To the best of our knowledge, no other approach considers this scenario. In order to provision for such conditions, the duration of the expected worst case period,  $\tau$  days, can be approximated at design-time, and the battery over-provisioned accordingly. For example, we might over-provision the battery with  $\sum^{\tau} (DC_e(d) \cdot P_{sys} \cdot 24 \text{ hours})$ , and let the duty-cycle be an exponentially decaying function for those days that are below some threshold  $E_t$ , i.e.,  $DC_e(d) = (DC_{min})^{\frac{d+1}{\tau}} \forall \{d | \hat{E}_{in}(d) < E_t\}$ , where  $DC_{min}$  is the minimum acceptable duty-cycle.

Using our approach to capacity planning from *Sec. 3.3*, enhanced with the above emergency provisioning, healthy discharge cycles during normal operation can be achieved, as the emergency store is only used in exceptional situations.

### 4.2.5 Algorithm Considerations

The proposed algorithm requires a constant amount of non-volatile memory to maintain  $W$  values of  $E_{real}(d)$ , which are necessary to compute  $\alpha$ . Furthermore, the system must be able to compute  $E_{in}(d) \forall d \in [1, 365]$  at runtime, or alternatively store 365 values representing  $E_{in}(d)$  as a look-up table. The computation time is linear with respect to the number of days for which  $E_{in}(d)$  is to be determined. Since our approach considers the source's long-term characteris-



tics, we are not concerned with sub-daily energy fluctuations. Hence, the sustainable performance level, *i.e.*, duty-cycle, for the entire day  $d + 1$  is computed only once at the end of day  $d$ . Finally, note that the capacity planning algorithm reviewed in *Sec. 3.3* is computed offline and relies only on one unknown parameter, *i.e.*,  $\Omega$ , which can be approximated easily. Similarly, for the dynamic performance scaling algorithm from *Sec. 4.1*, only the history window size  $W$  must be determined (see *Sec. 5.3.3*). In *Sec. 6* we demonstrate the algorithm’s feasibility for implementation in a real system.

Note that the algorithm presented in *Sec. 4.1* may be optimized. For example, rather than always starting with the modeled limits  $d_0$ ,  $d_1$ , and  $d_2$ , we may store the limits obtained on day  $d$  and use those as initial conditions on day  $d + 1$ . However, in the proof-of-concept implementation discussed herein, we are not concerned with the most efficient way to find the intersection of the two functions.

## 5 Experimental Evaluation

In this section we use extensive trace-driven simulations to compare the proposed dynamic power management scheme against several State-of-the-Art approaches. We show that the proposed algorithm achieves uninterrupted long-term operation, while outperforming the baseline approaches over a range of performance metrics. In *Sec. 6* we further exemplify the proposed technique’s performance using a real-world energy harvesting wireless sensing system.

### 5.1 Experimental Setup

#### 5.1.1 Baseline Algorithms

We compare our approach through simulation against State-of-the-Art (*SotA*) implementations of energy-predictive and battery-reactive approaches. Specifically, we implement the predictive duty-cycling scheme from [15] with two different energy predictors, *i.e.*, EWMA [15] and WCMA [18], and one reactive approach, *i.e.*, ENO-MAX [23]. Note that [18] only provides an energy prediction algorithm but does not discuss dynamic performance scaling, hence we use the scaling algorithm from [15] to compute the duty-cycle.

We have selected these particular algorithms for the following reasons. The technique in [15] achieves very good performance with minimal overhead, and is commonly used as a baseline for comparative analysis, *e.g.*, [18, 23]. It is also one of the few techniques that combines prediction and scheduling for solar harvesting systems. The technique in [18] has been shown to improve the prediction accuracy, but it has not been investigated if the improvement translates into increased system performance. Finally, the technique in [23] is a very well-performing representative of the class of battery-reactive approaches.

#### 5.1.2 Methodology and Simulation Input Data

To evaluate and compare the performance of the proposed solution, we simulate a solar energy harvesting system with the power management schemes introduced in *Sec. 5.1.1* and the trace data discussed in the following.

For the simulation input data, we resort to the National Solar Radiation Database<sup>1</sup> (NSRD) to obtain a twelve year dataset containing hourly solar radiation measurements at a

Table 1: Name, time-period, and location of NSRD<sup>1</sup> datasets used for evaluation of the proposed approach. Maximum, mean, minimum and variance of solar radiation are given in  $Wh$  for a panel with surface area  $A_{pv} = 15cm^2$ .

Name	Time Period	Lat [°]	Long [°]	Max	Avg	Min	Var
CA	1/98 – 12/09	34.05	-117.95	10.37	7.03	0.92	5.62
MI	1/98 – 12/09	42.05	-86.05	10.55	5.34	0.53	9.05
ON	1/98 – 12/09	48.05	-87.65	10.98	5.07	0.44	11.24

single observation point in California (CA), Michigan (MI), and Ontario (ON). The data for the first year of each dataset is used for calibration of the State-of-the-Art approach, and the remaining eleven years are used for simulation input data. The data traces (see *Table 1*) from the National Solar Radiation Database are given in  $Wh \cdot m^{-2}$  of solar energy incident on a flat surface with zero inclination. Hence, to account for smaller panel sizes, inefficiencies of individual components, and losses in energy storage during simulation, the data is conditioned with a typical efficiency of a midrange solar panel  $\eta_{pv} = 10\%$ , orientation angle  $\phi_p = 180^\circ$ , and inclination angle  $\theta_p = 0^\circ$ . We evaluated different panel sizes, but the results are comparable, hence we only show and discuss results for a panel with  $A_{pv} = 5cm^2$ . Finally, we consider battery charging and discharging efficiencies with  $\eta_{cc} = 0.9$  and  $\eta_{out} = 0.7$  which are reasonable efficiency factors.

#### 5.1.3 Simulation Details

The capacity planning technique from [15] is used to obtain the battery capacity  $B$  and supported power level  $P_{sys}$  at full performance (*i.e.*,  $DC = 100\%$ ) using one year of calibration data for each of the three datasets. We do the same with the capacity planning algorithm from *Sec. 3.3*, but do not provision for emergency situations (see *Sec. 4.2*). The results are shown in *Table 2*, and discussed in *Sec. 5.2*.

For each of the baseline implementations we use the authors’ recommended parameters, *i.e.*,  $K = 3$ ,  $D = 4$ ,  $\alpha = 0.3$  for WCMA [18], and  $\alpha = 0.5$  for EWMA [15]. For ENO-MAX [23], we use  $\alpha = 1/24$ , and  $\beta = 0.25$ . The authors suggest values between 0.25 and 0.75 for  $\beta$ , with lower values improving the duty-cycle stability at the cost of performance. We experimented with different values and noted negligible improvements in performance but noticeable increase in duty-cycle variance with increasing values for  $\beta$ . Finally, due to the hourly values given by the National Solar Radiation Database, we use  $N_w = 24$  instead of 48 daily update slots for EWMA, WCMA, and ENO-MAX. This results in a slight penalty in prediction accuracy, but significantly reduces computation complexity. Recall that our approach performs only one update per day, *i.e.*,  $N_w = 1$ .

We assume the battery to be fully charged at the start of the simulation, and simulate a low-power disconnect hysteresis of 60%, as commonly enforced by modern charge controllers [6]. This means that, if at any time the battery is fully depleted, the load will only be reconnected once the battery has been recharged to 60% of its capacity.

For the history window size used by the proposed algorithm, we assume  $W = 63$  days for all three datasets. The effects of this parameter are further discussed in *Sec. 5.3*.

<sup>1</sup>[http://tredc.nrel.gov/solar/old\\_data/nsrdb/1991-2010](http://tredc.nrel.gov/solar/old_data/nsrdb/1991-2010)

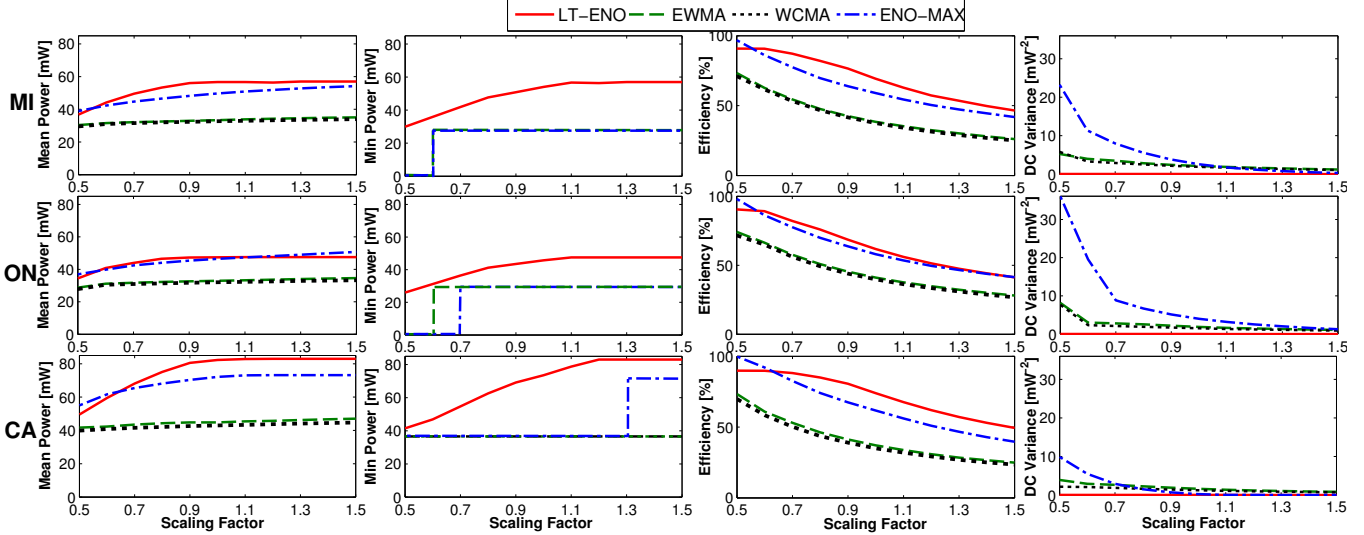


Figure 4: Mean power level, minimum power level, energy efficiency, and duty-cycle variance with energy input scaled (in increments of 10%) from 50% to 150% of original magnitude for each dataset and  $DC_{min} = 50\%$ .

Table 2: Battery capacities and supported power levels obtained with the *SotA* capacity planning approach and our proposed approach for the three datasets MI, ON, and CA (see Table 1), and a panel size  $A_{pv} = 5cm^2$ , and simulation results with fixed performance level, *i.e.*,  $DC(d) = 100\% \forall d$ .

	MI		ON		CA	
	<i>SotA</i>	LT-CP	<i>SotA</i>	LT-CP	<i>SotA</i>	LT-CP
$B$ [Wh]	42.51	93.6	61.48	98.61	56.96	88.96
$P_{sys}$ [mW]	55.73	57	56.75	47.5	73.17	83
Offline [%]	28.13	0	33.4	0	24.66	0
$P_{mean}$ [mW]	40.04	57	37.79	47.5	55.11	83
$P_{min}$ [mW]	0	57	0	47.5	0	83
DC Variance [ $\frac{1}{mW^2}$ ]	53.98	0	62.8	0	31.83	0

#### 5.1.4 Performance Metrics

Each of the algorithms are evaluated according to the following five performance metrics:

**Percent Time Offline.** For each experiment we report the percentage of the total simulation time during which the system was offline due to a depleted battery.

**Mean Power Level.** According to Sec. 3.1, the system utility is defined by the achievable duty-cycle. However, since the methods evaluated yield different sustainable power levels (see Table 2), we can not use the duty-cycle alone as a performance metric. Rather, we report the average power level achieved over all simulation time steps (including overriding zeros due to low-power disconnects), *i.e.*,  $P_{mean} = mean(DC(d)) \cdot P_{sys}$ .

**Minimum Power Level.** Achieving a minimum performance level can be crucial in certain application scenarios, *e.g.*, safety-critical systems. We therefore report the minimum power level that the evaluated approaches achieve.

**Duty-Cycle Variance.** We report the duty-cycle variance, normalized by the variance of  $E_{real}(d)$  over all simulation time steps.

**Energy Efficiency.** To compare the energy efficiency of the algorithms, we report the percentage of total energy that went unused because the battery was full.

## 5.2 Experimental Results

### 5.2.1 Capacity Planning

The State-of-the-Art (*SotA*) capacity planning algorithm discussed in [15] yields the required battery capacity  $B$  and sustainable power level  $P_{sys}$ , given an energy input trace representative of the conditions at the intended deployment site. Here, we investigate if  $P_{sys}$  can indeed be supported over long time periods by simulating the system equipped with a battery of capacity  $B$ , and running at a fixed, full performance power level  $P_{sys}$ , *i.e.*,  $DC(d) = 100\% \forall d$ , as obtained with the *SotA* capacity planning algorithm. The results in Table 2 show that the *SotA* approach does not always support the expected power level  $P_{sys}$ . The long-term energy neutral capacity planning (LT-CP) approach from Sec. 3, on the other hand, can sustain the expected performance level over the entire eleven years of simulation time, without relying on extensive trace data to calibrate the model.

### 5.2.2 Dynamic Power Management

According to [15], if the calibration data is representative of the actual conditions, the power level  $P_{sys}$  obtained with the *SotA* capacity planning technique should be supported at all times. However, the previous experiment showed that this may not always be the case, clearly demonstrating the need for dynamic power management. We thus evaluate and compare the dynamic power management approach proposed in this work against the performance of the power management techniques from [15] with EWMA and WCMA [18] predictors, and ENO-MAX [23]. For these algorithms we assume the power subsystem from *SotA* capacity planning [15], while our approach (LT-ENO) uses the capacity planning (LT-CP) reviewed in Sec. 3.3. We also analyze the scenario in which the baseline algorithms use batteries obtained with the LT-CP approach.



Table 3: Simulation results averaged over all simulation runs shown in Figure 4 for the three datasets, *i.e.*, MI, ON, CA, and the parameters listed in Table 2. *Note:* Static refers to capacity planning alone, *i.e.*, no DPM is used.

	Algorithm					
	Static	EWMA	WCMA	ENO-MAX	LT-ENO	
MI	Offline [%]	5.75	0.12	0.26	0.54	0
	$P_{mean}$ [mW]	53.71	33.22	32.42	48.63	52.77
	$P_{min}$ [mW]	36.27	25.33	25.33	25.33	49.49
	DC Var. $[\frac{1}{mW^2}]$	0.0001	2.5	2.29	5.35	0.0016
	Efficiency [%]	71.6	42.82	41.68	62.87	69.68
ON	Offline [%]	4.8	0.57	0.66	2.16	0
	$P_{mean}$ [mW]	45.21	32.65	31.66	45.45	45.25
	$P_{min}$ [mW]	30.22	25.79	25.79	23.22	41.97
	DC Var. $[\frac{1}{mW^2}]$	0.0005	2.47	2.06	8.25	0.001
	Efficiency [%]	64.6	45.3	43.82	62.58	64.4
CA	Offline [%]	6.29	0	0	0	0
	$P_{mean}$ [mW]	77.85	44.79	42.75	68.94	75.45
	$P_{min}$ [mW]	52.81	36.59	36.59	46.56	69.01
	DC Var. $[\frac{1}{mW^2}]$	0.0009	1.86	1.33	1.88	0.002
	Efficiency [%]	75.74	41.5	39.39	64.9	72.4

For simulation we fix the minimum acceptable duty-cycle at  $DC_{min} = 50\%$ . While this may seem like an unusually high duty-cycle, it is a reasonable lower bound considering that the power subsystem is designed such that a power level corresponding to  $DC = 100\%$  can be supported. We then simulate the different approaches with the energy input traces scaled from 50% to 150% to artificially cause model deviations. The results are shown in Figure 4, and Table 3 lists the performance results averaged over all simulation runs.

It is evident that the proposed approach (LT-ENO) outperforms the baseline algorithms in all respects, except for a few instances where the performance is comparable to that achieved by ENO-MAX. It is particularly noteworthy that the achieved mean power level is bounded closely by the minimum and maximum power levels respectively, illustrating a low duty-cycle variance. For the baseline algorithms, the achieved minimum power level is at most equal to the minimum acceptable power level, *i.e.*,  $P_{min} = P_{sys} \cdot DC_{min}$ . This means that for the baseline approaches, the minimum achieved power level follows the user defined minimum acceptable duty-cycle. With the proposed approach, however, the minimum achievable duty-cycle follows the long-term dynamics of the observed energy profile. Furthermore, considering long-term instead of short-term dynamics has a direct impact on duty-cycle variance. From Figure 4 and Table 3 it is evident that the duty-cycle variance is orders of magnitude lower than that obtained with any of the baseline algorithms. Achieving high duty-cycle stability over long time periods can be a strong requirement in a broad range of application scenarios, *e.g.*, [9, 22, 24].

From Figure 4 and Table 3 we further note that, while our approach achieves 100% availability in all simulation runs, the baseline algorithms suffer from depleted batteries for two of the three datasets. In the worst case, this results in system unavailability for up to 620 days (ENO-MAX with ON dataset scaled by 0.5). This behavior is expected since the baseline algorithms are battery agnostic and, as has been shown in Sec. 5.2.1, the power subsystem is underdimensioned. In order to perform a fair analysis, and determine if these algorithms could do better, we evaluate the

baseline algorithms with a power subsystem from Sec. 3.

The only significant difference to the results discussed above is with respect to the system’s availability, *i.e.*, the percentage time offline metric. The baseline algorithms now achieve 100% availability, *i.e.*, 0% offline, for all datasets. The little improvement in the other performance metrics is attributed to the battery agnostic nature of the baseline algorithms. In the case of the two predictive approaches, *i.e.*, EWMA and WCMA, an appropriately dimensioned battery only helps to overcome fundamental limitations of the approach, *i.e.*, short-term prediction. The reactive approach, *i.e.*, ENO-MAX, could benefit substantially from an appropriate battery if the setpoint required by this algorithm is computed dynamically according to an expected discharge profile that takes the battery capacity into consideration.

In this section, we have shown through simulation that our approach excels in all five performance metrics as defined in Sec. 5.1.4. The proposed dynamic power management scheme achieves 100% system availability in simulation with eleven years of trace data for different locations. We have shown that the minimum and mean expected performance level can be achieved even when there are deviations from the design-time model assumptions. Since our algorithm leverages the source’s long-term dynamics, an extremely low duty-cycle variance can be maintained while still achieving highly efficient energy usage.

### 5.2.3 Benefits of Dynamic Power Management

In the previous section we have shown that the proposed approach, which combines appropriate power subsystem capacity planning and a dynamic power management (DPM) scheme, yields considerable performance and reliability improvements when compared to the State-of-the-Art approaches proposed in literature. In this section we discuss the benefits of using our DPM algorithm over relying only on capacity planning.

For this purpose, we performed the same simulation discussed in the previous section, but set a static duty-cycle, as obtained from capacity planning. The results, averaged over all simulation runs are shown in the first column of Table 3. As is evident, the mean achievable duty-cycle without DPM support is approximately equal to the duty-cycle achieved by LT-ENO. However, if we consider the minimum achievable duty-cycle, the static approach performs significantly worse than LT-ENO. This is because the static approach can not adjust the performance level in response to deviations from the model, and experiences battery depletions when true conditions are below some percentage of the expected conditions.

For example, Figure 5 shows the minimum and mean duty-cycle achieved with the combination of capacity planning (for an expected duty-cycle  $DC_{sys} = 70\%$ ) and the proposed DPM algorithm over a range of scaling factors by which the energy input was scaled. The same is shown for capacity planning alone, *i.e.*, the static approach. Note that in region (a), *i.e.*, for scaling factors 0.5 to 0.9, the static approach was unable to sustain a non-zero minimum duty-cycle, inferring that the system suffered power outages. In this region, the DPM approach achieves a performance level roughly proportional to the expected duty-cycle scaled by the energy input scaling factor. In other words, the DPM ap-




Figure 5: Benefits of dynamic power management (DPM). Comparison of mean and minimum achievable duty-cycle with and without DPM.

proach trades off performance for ensuring continuous operation at an adjusted minimum expected duty-cycle that is dependent only on the energy input. This feature is a clear benefit for systems that require continuous operation.

Region (b) in *Figure 5*, *i.e.*, between 0.9 and 1.2 in this case, is a transition region where the static approach performs equivalently to DPM. Nevertheless, it is evident that the DPM enabled system is able to improve the performance level in response to increased energy availability. The lower bound of this region depends on the degree of overestimating true conditions, *i.e.*, an effect of capacity planning, while the upper bound of this region is dependent on the reactivity of the DPM algorithm, and therefore related to the history window size  $W$  (see *Sec. 5.3.3*). Finally, region (c) shows the full potential of DPM. The dynamic approach continues to adapt to the surplus energy and increases the performance level accordingly.

In summary, in this section we have shown that, unless the expected conditions at the intended deployment site can be very closely approximated, our DPM scheme provides two clear benefits. First, it allows reliable operation even when the expected conditions were significantly overestimated. Second, the algorithm can adapt to surplus energy, and safely increase the performance level accordingly.

### 5.3 Sensitivity Analysis

#### 5.3.1 Energy Profile Periodicity

The proposed dynamic power management scheme assumes a certain periodicity and sinusoidal behavior of the energy source. This is a valid assumption, since the tilt in the earth's axis of rotation will cause different incident angles depending on the annual solar cycle, which has a direct impact on the harvestable energy [7]. Despite assuming a stationary solar harvesting setup (*i.e.*, no tracking capabilities) it is nevertheless possible that the expected sinusoidal behavior fails to appear. For example, a natural, or man-made structure may shade the panel over the course of the year such that the typical peaks and troughs are obscured. Occurrences of these environmental effects are considered extenuating circumstances, and therefore not considered in this work. Nevertheless, in the following, we briefly investigate a similar effect due to the panel inclination angle.

#### 5.3.2 Panel Inclination and Orientation

The proposed approach builds upon an energy availability model with deployment specific input parameters. Here we briefly discuss the effects of orientation angle  $\phi_p$  and in-




Figure 6: Achievable maximum, mean, and minimum duty-cycle (top), and duty-cycle variance (bottom) with inclination angle  $\phi_p$  ranging from  $0^\circ$  to  $90^\circ$  for the CA dataset.

clination angle  $\theta_p$  of the solar panel. First, the effect of  $\phi_p$  has been considered in *Sec. 5.2*, where it was shown that the approach can handle significant deviations from the expected conditions. The inclination angle, however, changes the shape of the annual solar energy profile. Hence, to evaluate this effect, we simulate the system with the battery provisioned as before, *i.e.*,  $\theta_p = 0^\circ$  for the CA dataset, but vary the panel inclination angle  $\phi_p$  from  $0^\circ$  to  $90^\circ$  for the simulation. *Figure 6* shows the maximum, mean, and minimum achievable duty-cycle for this modified CA dataset. Note that the minimum allowable duty-cycle is fixed at 1%. The results clearly show that the proposed dynamic power management scheme can adapt to unexpected energy profiles, while maintaining very low duty-cycle variance.

#### 5.3.3 History Window Size

When applying a large history window size  $W$ , the scaling factor  $\alpha$  contains information about environmental conditions  $W$  days in the past. Large  $W$  values reduce the model's reactivity to significant variations in the present energy profile, which has the potential to threaten uninterrupted operation. On the other hand, a short history window enables reacting to short-term variations of the source, but at the cost of increased duty-cycle variance.

In order to find a suitable trade-off between achievable performance level and duty-cycle stability, we evaluate the adjusted model's approximation accuracy with different values for  $W$ . We define the performance metric given in (10), which considers the model's approximation accuracy through Mean Absolute Percentage Error (MAPE) [1], scaled by the variance of  $\alpha$  over a time period of  $N$  days.

$$\sigma = \frac{\text{var}(\alpha)}{N} \sum_{i=1}^N \left( \left| \frac{E_{\text{real}}(d_i) - \hat{E}_m(d_i)}{E_{\text{real}}(d_i)} \right| \right) \quad (10)$$

The result for  $W \in [7, 140]$  days, and  $N = 365$  days is illustrated in *Figure 7*, which shows  $\sigma$  for the three different datasets, normalized by the respective maximum value of  $\sigma$ . As expected, with increasing  $W$ , the performance metric  $\sigma$  approaches a value, past which there is diminishing improvement in approximation accuracy or stability. Intuitively, the optimal value for  $W$  is likely dependent on the source characteristics, *i.e.*, the variability of the energy profile, and on the length of the deficit period that the battery must be able to bridge. Based on the results in *Figure 7*, a history window of  $W = 63$  days for all three datasets is considered a suitable parameterization.




Figure 7: Evaluation of history window size  $W$ , normalized to maximum value of  $\sigma$  for the respective dataset.

## 6 Case Study

In this section we demonstrate the feasibility for realization of the theoretical models formulated in *Sec. 3* and *4* in a real-world system. In the context of a case study, we quantify the benefit of the proposed dynamic power management scheme in terms of increased application utility.

### 6.1 System Description

#### 6.1.1 Motivation

The X-SENSE project’s [3] goal is to (i) apply WSN technology to enable geoscientific characterization and quantification of cryosphere phenomena, and their transient responses to climate change, and (ii) investigate the feasibility of WSN technology for an early-warning system against destructive events triggered by these phenomena. An integral part of the geoscientific aspects of the project is the ability to track movement on the order of a few centimeters per month. This requires a system that can reliably, and autonomously provide positioning with sub-centimeter accuracy over long time periods, *i.e.*, multiple years.

For this reason, a number of in-situ evaluation, and experimentation platforms enhanced with commercially available single-frequency GPS receivers are installed at the deployment site (see *Figure 8b*). The raw L1 GPS data collected is differentially post-processed [5] to yield the positioning accuracy required by the application [24]. Long sampling intervals required for an acceptable solution accuracy (see *Sec. 6.1.3*), coupled with the GPS receivers’ high power demands [14], and the need for unattended operation over multiple years necessitate that an energy harvesting system is employed to power the experimentation platform.

The project requirements and characteristics of the experimentation platform present an ideal application scenario for the validation of assumptions and models in *Sec. 3* and *4*, and the evaluation of the dynamic power management scheme under real-world conditions.

#### 6.1.2 System Architecture

The system architecture is illustrated in *Figure 8a* and a picture of a deployed X-SENSE system is shown in *Figure 8b*. It consists of the solar energy harvesting and energy storage subsystems, collectively referred to as power subsystem, and the load, *i.e.*, wireless sensing platform, to be supported.

**Wireless Sensing Platform.** To satisfy the project requirements from *Sec. 6.1.1*, we leverage a custom-built, feature-rich hardware platform together with an extensible middleware, introduced in [4]. With this platform we trade off system flexibility, system observability, and accessibility for rel-

Table 4: Name, time-period, coordinates, and solar panel orientation ( $\phi_p$ ) and inclination ( $\theta_p$ ) angles of the deployed systems used in the case study.

Name	Time Period	Lat [°]	Long [°]	$\phi_p$ [°]	$\theta_p$ [°]
DH	02/01/12 – 03/22/14	46.1235531	7.82126695	195	57.5
GG	03/16/12 – 03/22/14	46.0901923	7.81339546	210	65

atively high power dissipation. Since the high load could not be supported over extended periods of time with batteries alone, the system is powered by the solar energy harvesting system discussed below.

Although proven to be a very powerful and flexible tool for in-situ experimentation, the evaluation platform has one major limitation: it can not measure the current generated by the panel  $I_{pv}$ , and only monitor the system input voltage  $V_{sys}$  and current drawn  $I_{sys}$  (see *Figure 8a*). However, as discussed in *Sec. 4*, an approximation of the generated energy  $E_{real}(d)$ , which can be derived from  $I_{pv}$ , is necessary for the DPM algorithm to function. *Sec. 6.1.4* explains how we circumvent the lack of appropriate hardware support.

**Power Subsystem.** As illustrated in *Figure 8a*, and mentioned in *Sec. 3.1*, we assume a harvest-store-use architecture [20], which is enforced by the employed charge controller. This means that the energy to operate the system is always supplied by the battery, even when the panel generates surplus energy, *e.g.*, when the battery is full.

We consider two configurations that are identical with respect to technology parameters (see *Sec. 3*): a 30 Watt mono-crystalline solar panel (cleversolar CS-30) with a solar cell area of  $0.1725m^2$ , a SunSaver SS-6L PWM charge controller, a Lifeline AGM battery with a nominal capacity of 54Ah, and the wireless sensing system described above with a power dissipation of approximately 6 Watt at full performance. *Table 4* lists the relevant deployment parameters that differ between the two configurations.

Note that, according to our model in *Sec. 3* and *4*, neither of the two configurations is ideal with respect to the dimensioning of the energy harvesting power subsystem. The panel and battery are poorly matched: the battery is under-provisioned for the given load (the deployed battery is only dimensioned for 10 days of operation at a duty-cycle of 30%, and hence not large enough to bridge extended periods of snow cover during winter), and the panel is over-provisioned for the battery employed (resulting in available energy being wasted during summer time). Nevertheless, using the capacity planning algorithm from *Sec. 3* and assuming  $\Omega = 0.6$ , we find that this harvesting configuration should be able to support 7.5 hours (*i.e.*,  $DC = 31.25\%$ ) of daily system operation for DH, and 6.7 hours for GG, as long as any one period without harvesting opportunities does not exceed 10 days. In order to reduce the likelihood of a low-power disconnect due to an under-provisioned battery, we set the history window size  $W = 10$  days (see *Sec. 4.1* and *5.3.3*) to match the period that can be bridged by the battery.

#### 6.1.3 System Utility

In *Sec. 3.1* we defined the system utility to be directly proportional to the duty-cycle that the system can achieve. Here we refine it according to the aforementioned case study.

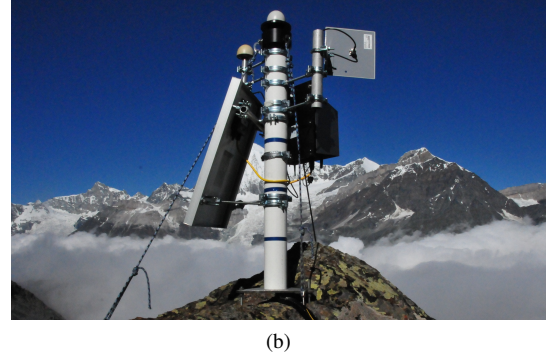
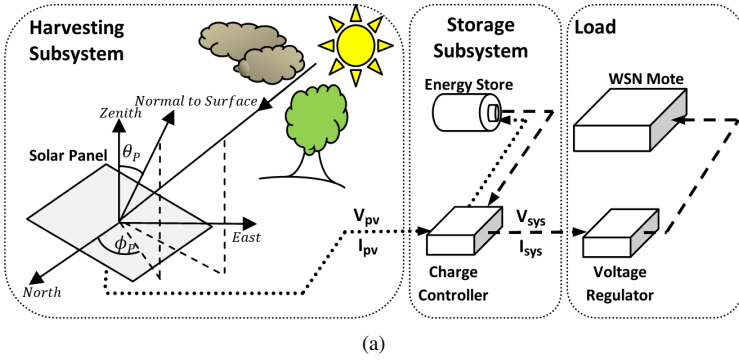


Figure 8: (a) System block diagram partitioned into harvesting subsystem (solar panel), storage subsystem (battery and charge controller) and the load to be supported (wireless sensing system). The flow of generated energy is indicated with dotted lines, while the dashed lines represent the energy consumed. (b) Picture of an X-SENSE energy harvesting wireless sensing system installed at the high-alpine deployment site.

As introduced in *Sec. 6.1.1*, the application scenario [24] relies on a differential GPS processing algorithm, whose acceptable error performance requires periodic sampling of the GPS receiver ( $\mu$ -blox LEA-6T) over at least two consecutive hours per day. This is illustrated in the top graph in *Figure 9*, which shows the processing algorithm's error performance as a function of the measurement duty-cycle. From the figure, it is evident that the minimum acceptable error, *i.e.*,  $e(DC) \leq 8\text{mm}$ , requires at least two hours of continuous sampling [5], above which the error decreases exponentially. Note that the error is undefined for duty-cycles lower than approximately 8%, as this is below the minimum required by the processing algorithm.

The application under consideration directly benefits from increased temporal resolution, *e.g.*, to characterize sub-daily process variations, hence, sampling over longer time intervals increases the system utility. Therefore, we define the system utility, as shown on the bottom graph in *Figure 9*, to be  $U(DC) = 1 - e_{norm}(DC)$ , where  $e_{norm}(DC)$  represents the normalized error performance  $e(DC)$  shifted by an offset to reach the maximum utility at a duty-cycle of 100%. This is a realistic definition of system utility for many application scenarios, *e.g.*, [9, 11].

#### 6.1.4 Implementation Details

In the following we discuss implementation aspects that are relevant for the particular wireless sensing system under consideration. By discussing two adaptations due to technical limitations of the system, we demonstrate the feasibility for implementation even in systems that are not specifically designed with our approach in mind.

**Circumventing Limited Measurement Support.** As already mentioned, for a number of reasons, the hardware platform is not designed to provide measurements of the energy generated by the solar panel, or the energy actually flowing into the battery. However, the dynamic power management scheme introduced in this work relies on an approximation of the energy harvested over a given day. In order to approximate the harvested energy without appropriate hardware support, we leverage a recently proposed battery State-of-Charge (SOC) algorithm [6]. At an absolute min-




Figure 9: Error performance of processing algorithm versus duty-cycle (top), and derived system utility versus duty-cycle (bottom). Note that the error performance is undefined for a duty-cycle lower than 8%.

imum, this algorithm requires measurements of the system input voltage  $V_{sys}$  and current discharge rate  $I_{sys}$  (*i.e.*, electric current drawn by the system) to approximate the battery fill level, *i.e.*, SOC. The State-of-Charge indication and the length of daily charging cycles provided by this algorithm are used to obtain an approximation of  $E_{real}(d)$  as follows.

In the case when the power subsystem is not optimally provisioned, the State-of-Charge alone is not sufficient to approximate the energy generated by the panel. For instance, if the battery is full, any surplus energy generated by the panel will be dissipated by the charge controller. This condition will not be visible to the State-of-Charge algorithm used herein. Therefore, we assume  $T_c(d)$  to be the duration of the daily charging cycle, *i.e.*, the duration over which the panel generated enough power to keep the charge controller in charging mode, as given by the State-of-Charge algorithm. Then, given the panel's maximum power rating  $P_{pv}$ , we can approximate the maximum energy  $\hat{H}_{max}$  that can be harvested on a given calendar day  $d$  as  $\hat{H}_{max}(d) = T_c(d) \cdot P_{pv}$ .

Since we are interested in the energy actually generated by the panel, and not its theoretical daily maximum, we scale  $\hat{H}_{max}(d)$  by a factor  $\zeta$  to approximate  $E_{real}(d)$ . The scaling factor  $\zeta$  accounts for the fact that, with the given configuration, a certain fraction of the energy generated by the panel is wasted. We approximate this scaling factor as shown in (11) by considering only the days on which the State-of-Charge



(SOC) approximation is below 80%.

$$\zeta_d = \text{mean} \left( \zeta_{d-1} + \frac{G(d)}{\widehat{H}_{\max}(d)} \right), \text{ if } \text{SOC}(d) \leq 80\%, \forall d \quad (11)$$

The quantity  $G(d)$  is the daily energy generation approximated using measurement of the system’s current drain, and the State-of-Charge approximation, as given in (12). Note,  $V_{bat}$  is the operating voltage, *i.e.*,  $V_{bat} = 12V$  in our case, and  $B$  is the battery capacity, *i.e.*,  $B = \eta_{out} \cdot B_{nom}$ .

$$G(d) = \frac{\widehat{E}_{out}(d)}{\eta_{in} \cdot \eta_{out}} + \frac{(\text{SOC}(d) - \text{SOC}(d-1)) \cdot B \cdot V_{bat}}{\eta_{in}} \quad (12)$$

**DPM Algorithm Modification.** The dynamic power management (DPM) algorithm discussed in *Sec. 4.1* was designed for enabling uninterrupted long-term operation as the primary goal. The algorithm assumes an appropriately provisioned battery and solar panel, as obtained with the capacity planning from *Sec. 3.3*. However, as already discussed, the configuration under consideration is suboptimal, hence we make the following adjustment.

Since the panel is over-provisioned, it will likely be able to generate significantly more energy during periods of surplus than what can be stored in the battery, and used by the system (see *Sec. 3.2*). Hence, during periods of surplus (*i.e.*, on the interval  $[d_0, d_1)$ ), we scale the duty-cycle computed by the DPM algorithm by a factor  $\psi = 2$ . This value has been selected based on experiments, which suggest that the exact value for  $\psi$  depends on the ratio of energy consumption and generation, but this must be investigated further. In order to reduce excessive usage of the energy stored in the battery, we introduce a guard time  $T_g$  of 30 days. This means that the modification will only be used on day  $d$ , if and only if the following inequality holds:  $d_0 + T_g < d < d_1 - T_g$ .

## 6.2 Performance Evaluation

This section presents and discusses the results obtained from running the proposed dynamic power management algorithm on the system introduced in *Sec. 6.1*. In order to verify the simulation framework’s applicability for investigating different algorithm parameterizations, we first briefly discuss our validation methodology. We recorded traces<sup>2</sup> of battery State-of-Charge, harvested energy, and load of the wireless sensor platform described in *Sec. 6.1.2*. We then set the parameters of our simulation framework to represent this system and compared the recorded State-of-Charge trace to the output of the simulation framework. We observed a Mean Absolute Percentage Error [1] between measured and simulated battery State-of-Charge of 5.45% over the entire trace. We therefore conclude that simulation with other algorithmic parameterizations will exhibit similar low deviations.

**Results.** The graphs on top in *Figure 10* show the daily energy  $E_{real}(d)$  together with the modeled energy expectation  $E_{in}(d)$  and the modeled energy consumption  $E_{out}(d)$  for each system over almost 700 days. Recall that  $E_{in}(d)$  is derived from the astronomical model at design-time, and, depending on its parameterization can be an optimistic or pessimistic

approximation. The bottom graphs show the static duty-cycle expected ( $Static_{exp}$ ) based on model assumptions, and the dynamically computed duty-cycle  $LT - ENO^*$ , which incorporates the modification discussed in *Sec. 6.1.4*. For reference, we also show the computed duty-cycle if this modification were to be disabled (referred to as  $LT - ENO$ ), and the static duty-cycle  $Static_{max}$  that may actually be supported by the respective energy harvesting configuration when no dynamic power management is employed. Note that for the latter, we assume that perfect knowledge of the true energy conditions are available at design time. Finally, to investigate the performance boundaries of our algorithm, the graphs show the duty-cycle achieved by a clairvoyant energy-prediction algorithm referred to as *Clairvoyant*. Note that, while this algorithm has perfect knowledge of a time window of 30 days into the future, it is *not* necessarily a perfect scheduler.

From the top graphs in *Figure 10*, we note that the expected harvesting opportunities were significantly overestimated with the selection of the environmental parameter  $\Omega = 0.6$  (see *Sec. 3.2*). A more appropriate, and safer parameterization would be  $\Omega = 0.83$  for DH, and  $\Omega = 0.8$  for GG. This results in a supported duty-cycle  $Static_{max} = 12.9\%$  for DH, and  $Static_{max} = 10.8\%$  for GG, assuming all other parameters are unchanged. From the bottom graphs it is evident that  $LT - ENO$  stays between the expected duty-cycle based on model assumptions  $Static_{exp}$  and the supported static duty-cycle  $Static_{max}$  for most of the time. This clearly shows the proposed algorithm’s ability to dynamically adapt the performance level in response to deviations from the modeled expectations, without risking battery depletion. As is evident from *Table 5*, by using  $LT - ENO$ , power outages due to overestimating actual conditions, as experienced by the static approach, *i.e.*,  $Static_{exp}$  can be eliminated while simultaneously improving both system utility and energy efficiency.

The results for GG in *Table 5* show that, although the system achieves zero downtime, the minimum achieved duty-cycle is below the minimum duty-cycle required by the end-user application. The resulting periods with zero utility (see *Sec. 6.1.3*) are due to time periods with no harvesting opportunities that significantly exceed the time period the battery can bridge. As discussed in *Sec. 4.2.4*, zero input conditions trigger an emergency mechanism, which overrides the duty-cycle computed by our algorithm due to lack of energy input around days 105 and 690 on the bottom graph in *Figure 10b*. Although experiencing days with zero energy input for DH as well, these instances are not as prominent in the dataset.

Next we consider the modified approach represented by  $LT - ENO^*$ , which differs from  $LT - ENO$  only during periods of energy surplus, *i.e.*, summer. The effect of the modification discussed in *Sec. 6.1.4* is clearly visible by the sudden step in duty-cycle around days 300, 400, and 670 for DH in *Figure 10a*, and days 290, 430, and 650 for GG shown in *Figure 10b*. We note that  $LT - ENO^*$  approaches the performance of the clairvoyant algorithm. In both cases, the duty-cycle achieved by the clairvoyant algorithm tightly bounds that of  $LT - ENO^*$  during periods of energy deficit, *i.e.*, winter. On average over the entire period,  $LT - ENO^*$  achieves 76.26% of the performance of the clairvoyant algorithm for DH and 68.15% for GG.

<sup>2</sup>All datasets used in this work are available at <http://data.permasense.ch>

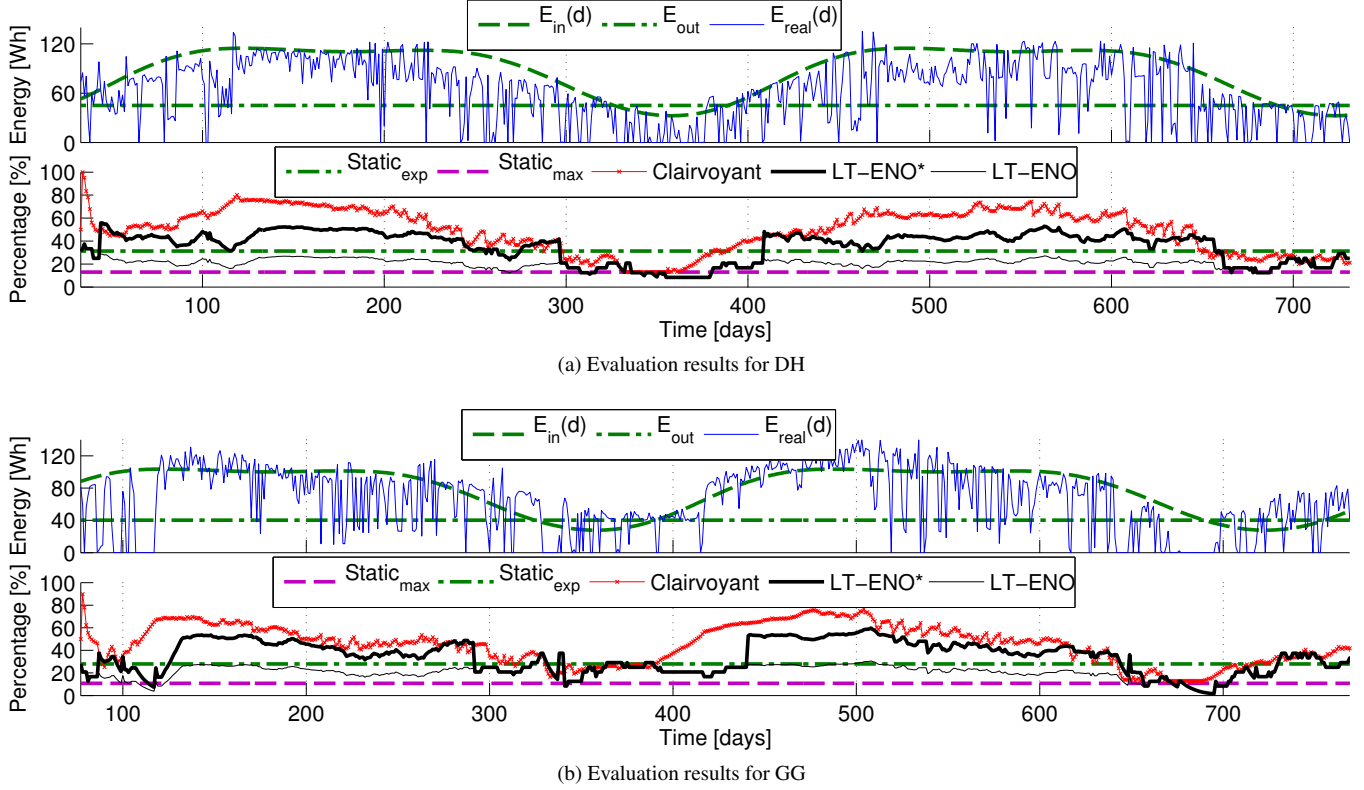


Figure 10: Uninterrupted operation over 700 days for two energy harvesting systems using the LT-ENO DPM approach. Note that  $E_{in}(d)$  refers to the offline model, which, in this case, consistently overestimated the true conditions. In general,  $E_{in}(d)$  can be an optimistic or pessimistic approximation.

Table 5: Percent Offline, mean ( $\overline{DC}$ ) and minimum ( $\lfloor DC \rfloor$ ) duty-cycle, duty-cycle variance ( $\sigma^2$ ), mean utility ( $\overline{U}$ ), minimum utility ( $\lfloor U \rfloor$ ), and energy efficiency ( $\eta$ ) for DH and GG shown in Figure 10.

Note:  $LT-ENO$  is the original DPM algorithm,  $LT-ENO^*$  is the DPM algorithm modified according to Sec. 6.1.4,  $Static_{exp}$  and  $Static_{max}$  are expected static duty-cycle according to model assumptions.

	Power Management Algorithm	Offline [%]	$\overline{DC}$ [%]	$\lfloor DC \rfloor$ [%]	$\sigma^2$	$\overline{U}$ [%]	$\lfloor U \rfloor$ [%]	$\eta$ [%]
DH	$LT-ENO^*$	0	35.8	8.3	0.0167	81.2	22.4	80.3
	$LT-ENO$	0	20.6	8.3	0.0021	70.3	22.4	44.5
	$Static_{exp}$	7.9	28.8	0	0.0004	68.9	0	62.5
	$Static_{max}$	0	12.9	12.9	0	49.0	49.0	26.7
GG	$LT-ENO^*$	0	34.79	1.43	0.0207	79.0	0	75.0
	$LT-ENO$	0	21.7	1.43	0.0044	69.1	0	46.5
	$Static_{exp}$	8.0	25.7	0	0.0004	69.2	0	54.6
	$Static_{max}$	0	10.8	10.8	0	38.6	38.6	21.7

The performance improvements due to efficient use of the power subsystem are ultimately expected to translate into high system utility for the end-user. Figure 11 shows the histogram of system utility (defined in Sec. 6.1.3) achieved by the four different approaches. The static approach with  $Static_{max}$  achieves a constant, but clearly the lowest utility. As is to be expected, the overly optimistic  $Static_{exp}$  increases the overall utility at the cost of a few days with zero utility. Note that the system achieved 100% uptime, but not always a sampling time of at least 2 hours, thus resulting in zero utility. As discussed in Sec. 5.2.3, the two dynamic approaches, i.e.,  $LT-ENO$  and  $LT-ENO^*$ , trade-off performance for




Figure 11: Histogram of utility for the static ( $Static_{max}$  and  $Static_{exp}$ ) and dynamic approaches ( $LT-ENO$  and  $LT-ENO^*$ ).

ensuring continuous operation. Significantly higher utility can be achieved as a result of the system's ability to react to deviations from expected conditions and efficiently adjusting the performance to safe levels.

In summary, we have shown that the proposed approach to enabling long-term uninterrupted operation of solar energy harvesting systems is very applicable even to systems with limited hardware support. Using a real deployed sensing system, we have demonstrated that the proposed approach results in significant improvements in system utility.

## 7 Conclusions

In this work we have demonstrated that appropriate design-time considerations, together with a novel dynamic



power management scheme can indeed enable energy neutral operation of solar energy harvesting systems over time periods on the order of multiple years. The proposed dynamic power management scheme leverages an astronomical energy availability model that is also used to dimension the energy harvesting power subsystem.

Rather than considering the energy source's short-term fluctuations, our approach uses the source's known long-term tendencies to compute the sustainable duty-cycle. This allows the system to fully, but safely, leverage the power subsystem and so achieve stable performance over long time periods without incurring downtime. Additionally, the systematic end-to-end solution enables efficient use of the power subsystem, resulting in major savings in terms of system cost and physical form factor.

When compared to the State-of-the-Art implementations, the proposed approach achieves a reduction in duty-cycle variance by up to three orders of magnitude without impeding the achievable performance level. Trace-driven simulation with eleven years of real-world data showed that the achieved minimum, and mean duty-cycle improve upon existing techniques by up to 195% and 177% respectively. Finally, using a concrete case study, we have shown that the proposed approach can significantly improve system utility, while exhibiting robustness against variations in the observed energy profile, irrespective of the source of model deviations, *i.e.*, environmental and deployment variations.

## 8 Acknowledgments

This work was scientifically evaluated by the SNSF and financed by the Swiss Confederation and by Nano-Tera.ch. The authors would like to thank the anonymous reviewers and Guoliang Xing for their valuable feedback.

## 9 References

- [1] J. Armstrong and F. Collopy. Error measures for generalizing about forecasting methods: Empirical comparisons. *International Journal of Forecasting*, 8(1):69–80, 1992.
- [2] H. J. Bergveld. *Battery management systems: design by modelling*. PhD thesis, Enschede, June 2001.
- [3] J. Beutel et al. X-Sense: Sensing in extreme environments. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2011*, pages 1–6. IEEE, 2011.
- [4] B. Buchli et al. Demo abstract: Feature-rich platform for WSN design space exploration. In *Proceedings of the 10th international conference on Information Processing in Sensor Networks*, pages 115–116, 2011.
- [5] B. Buchli et al. GPS-equipped wireless sensor network node for high-accuracy positioning applications. In *Wireless Sensor Networks*, pages 179–195. Springer, 2012.
- [6] B. Buchli et al. Battery state-of-charge approximation for energy harvesting embedded systems. In *Wireless Sensor Networks*, pages 179–196. Springer, 2013.
- [7] B. Buchli et al. Towards Enabling Uninterrupted Long-Term Operation of Energy Harvesting Embedded Systems. In *Proc. of 11th European Conference on Wireless Sensor Networks (EWSN 2014)*, Oxford, UK, Feb 2014. Springer Link, Lecture Notes on Computer Science.
- [8] M. Buzzi. *Challenges in operational numerical weather prediction at high resolution in complex terrain*. PhD thesis, Diss., Eidgenössische Technische Hochschule ETH Zürich, Nr. 17714, 2008.
- [9] M. Chang and P. Bonnet. Meeting ecologists' requirements with adaptive data acquisition. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 141–154. ACM, 2010.
- [10] P. Corke et al. Long-duration solar-powered wireless sensor networks. In *Proceedings of the 4th workshop on Embedded networked sensors, EmNets '07*, pages 33–37, New York, NY, USA, 2007. ACM.
- [11] V. Dyo et al. Evolution and sustainability of a wildlife monitoring sensor network. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 127–140. ACM, 2010.
- [12] T. He et al. Achieving Real-Time Target Tracking Using Wireless Sensor Networks. In *Real-Time and Embedded Technology and Applications Symposium. Proc. of the 12th IEEE*, pages 37–48, 2006.
- [13] D. Heinemann et al. Forecasting of solar radiation. *Solar energy resource management for electricity generation from local level to global scale*. Nova Science Publishers, New York, 2006.
- [14] R. Jurdak et al. Adaptive GPS duty cycling and radio ranging for energy-efficient localization. In *Proc. of the 8th Conference on Embedded Networked Sensor Systems*, pages 57–70. ACM, 2010.
- [15] A. Kansal et al. Power management in energy harvesting sensor networks. *ACM Transactions on Embedded Computing Systems*, 6(4):32, 2007.
- [16] T. N. Le et al. Power Manager with PID controller in Energy Harvesting Wireless Sensor Networks. In *Green Computing and Communications (GreenCom), 2012 IEEE International Conference on*, pages 668–670. IEEE, 2012.
- [17] J. Lu and K. Whitehouse. SunCast: fine-grained prediction of natural sunlight levels for improved daylight harvesting. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, pages 245–256. ACM, 2012.
- [18] J. R. Piorno et al. Prediction and management in energy harvested wireless sensor nodes. In *Wireless Communication, Vehicular Technology, Information Theory and Aerospace & Electronic Systems Technology, 2009. Wireless VITAE 2009. 1st International Conference on*, pages 6–10. IEEE, 2009.
- [19] N. Sharma et al. Cloudy computing: Leveraging weather forecasts in energy harvesting sensor systems. In *Sensor Mesh and Ad Hoc Communications and Networks (SECON), 2010 7th Annual IEEE Communications Society Conference on*, pages 1–9. IEEE, 2010.
- [20] S. Sudevalayam and P. Kulkarni. Energy harvesting sensor nodes: Survey and implications. *Communications Surveys & Tutorials, IEEE*, 13(3):443–461, 2011.
- [21] J. Taneja et al. Design, Modeling, and Capacity Planning for Micro-solar Power Sensor Networks. In *Proc. of the 7th international conference on Information processing in sensor networks, IPSN '08*, pages 407–418, Washington, DC, USA, 2008. IEEE Computer Society.
- [22] I. Vasilescu et al. Data Collection, Storage, and Retrieval with an Underwater Sensor Network. In *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems, SenSys '05*, pages 154–165, New York, NY, USA, 2005. ACM.
- [23] C. M. Vigorito et al. Adaptive control of duty cycling in energy-harvesting wireless sensor networks. In *Sensor, Mesh and Ad Hoc Communications and Networks, 2007. SECON'07. 4th Annual IEEE Communications Society Conference on*, pages 21–30. IEEE, 2007.
- [24] V. Wirz et al. Temporal Characteristics of Different Cryosphere-related Slope Movements in High Mountains: GPS Measurements and Analysis. In *Abstract Volume 9th Swiss Geoscience Meeting*, volume 9, page 282. Platform Geosciences, Swiss Academy of Science, Nov 2011.