

Towards Enabling Uninterrupted Long-Term Operation of Solar Energy Harvesting Embedded Systems

Bernhard Buchli, Felix Sutton, Jan Beutel, Lothar Thiele

Computer Engineering and Networks Laboratory, ETH Zurich, Zurich, Switzerland
{bbuchli, fsutton, beutel, thiele}@tik.ee.ethz.ch

Abstract In this work we describe a systematic approach to power subsystem capacity planning for solar energy harvesting embedded systems, such that uninterrupted, long-term (*i.e.*, multiple years) operation at a predefined performance level may be achieved. We propose a power subsystem capacity planning algorithm based on a modified astronomical model to approximate the harvestable energy and compute the required battery capacity for a given load and harvesting setup. The energy availability model takes as input the deployment site's latitude, the panel orientation and inclination angles, and an indication of expected meteorological and environmental conditions. We validate the model's ability to predict the harvestable energy with power measurements of a solar panel. Through simulation with 10 years of solar traces from three different geographical locations and four harvesting setups, we demonstrate that our approach achieves 100% availability at up to 53% smaller batteries when compared to the state-of-the-art.

Keywords: Wireless sensor networks; energy harvesting; modeling; experimentation.

1 Introduction

Advances in miniaturization and low-power design of electronic devices have allowed Wireless Sensor Networks (WSNs) to reach a state at which they represent a feasible option for continuous observation of various processes, *e.g.*, industrial, agricultural, or environmental monitoring. Due to remote and inaccessible deployment sites, and the lack of power sources, the motes comprising a sensor network are usually battery powered devices [1]. However, the finite energy store imposed by non-rechargeable batteries severely limits the achievable performance level of application scenarios with increased energy demands, *e.g.*, [2]. Ambient energy harvesting, particularly in the form of solar energy harvesting [3, 4], has thus attracted much attention as a promising solution for enabling perpetual system operation.

The ultimate goal of an energy harvesting system is to enable uninterrupted long-term (*i.e.*, on the order of multiple years) operation at a defined minimum performance level (*i.e.*, sustained duty-cycle). However, simply enhancing an embedded system with energy harvesting capabilities may not suffice to achieve this goal [5, 6]. This is because solar energy harvesting opportunities depend both on static and dynamic factors, *e.g.*, efficiencies of the solar panel and energy storage element, solar panel installation parameters [5], and time-varying meteorological conditions, and transient local obstructions. Despite the many design examples that can be found in literature describing solar energy harvesting systems specifically designed for long-term WSN applications,

the problem of how to systematically sizing the energy store and solar panel capacities, such that uninterrupted, long-term operation at a defined performance level can be achieved, remains an open task. The work presented in this paper addresses this issue and provides a tool for design space exploration and identification of the smallest combination of battery and solar panel for a particular application.

The performance level achievable by an embedded computing system is ultimately limited by the available energy. Thus, the battery capacity must be large enough to cover the longest period without harvesting opportunities, and the panel must be able to generate sufficient amounts of energy to replenish the battery within an acceptable time frame. However, due to highly dynamic meteorological conditions, it may be difficult to define either of these periods without resorting to pessimistic assumptions [7]. In fact, literature review shows that most efforts attempt to mitigate the effects of improperly designed power subsystems with dynamic load scaling [8], which are based on expected harvesting opportunities predicted by an energy prediction scheme, *e.g.*, [9, 10].

While runtime energy awareness and reactivity is important for achieving improved system utility and energy efficiency, we argue that it can not replace appropriate design-time capacity planning of the power subsystem. Even a perfect, yet fictional energy predictor could achieve continuous operation only if the *entire* power subsystem is provisioned to support the given load. Hence, for application scenarios that require uninterrupted long-term operation at a pre-defined minimum performance level, the aforementioned approaches alone are unsatisfactory because they do not consider the effects of the power subsystem capacities.

To this end we propose a design-time, *i.e.*, offline power subsystem capacity planning algorithm for solar energy harvesting systems. The approach considers seasonal variations of the energy source, the sun, to approximate the harvestable energy and compute the required battery capacity given a panel size and deployment location information. Assuming that the modeled conditions reflect actual conditions, the power subsystem capacity obtained with this approach enables uninterrupted operation at a defined performance level, subject only to hardware failure, or environmental phenomena with a long-term effect on harvesting opportunities, *e.g.*, volcanic eruption.

The contributions of this work are summarized as follows. First, we propose a power subsystem capacity planning algorithm based on a modified astronomical model [11]. Second, we validate the modified model's ability to approximate the harvestable energy with power measurements of a solar panel. Thirdly, we evaluate the capacity planning approach through simulation, and show that it yields smaller batteries than the state-of-the-art (*SotA*) capacity planning algorithm [12]. We further show that our approach outperforms the *SotA* by achieving 100% availability over ten years for three different test datasets, while requiring up to 53% smaller batteries. In contrast to *SotA*, our approach does not rely on detailed power traces or calibration data, but only requires a crude estimate of average meteorological conditions at the intended deployment site.

Sec. 2 briefly reviews related work. *Sec. 3* introduces the high-level concept of the power subsystem capacity planning approach. *Sec. 4* provides a detailed discussion of the harvesting conditioned energy availability model that forms the basis of the capacity planning algorithm presented in *Sec. 5*. In *Sec. 6*, we evaluate the proposed method's ability to achieve uninterrupted long-term operation, and its sensitivity to parameter

variations. Finally, *Sec. 7* concludes this work with a summary and brief discussion of future work.

2 Related Work

Many design examples of energy harvesting systems can be found in literature, *e.g.*, ZebraNet [13], Heliomote [14], Ambimax [15], Fleck [16], Rivermote [17], *etc.* However, they fail to provide systematic approaches for power subsystem capacity planning, and instead present anecdotal, application specific design choices that are based on simplified assumptions. Since the realization that energy harvesting is not necessarily sufficient to guarantee uninterrupted operation [7], efforts have primarily focused on mitigating the impacts of an inappropriately provisioned power subsystem with energy prediction schemes, *e.g.*, [18–20], and dynamic load scheduling based on short-term predictions [8–10, 21]. Three notable exceptions are [7, 12, 22], which are briefly discussed in the following.

In [12] an analytical model for long-term sustainable operation is presented. The authors consider battery capacity planning based on a representative power profile, inferring that the panel size is fixed. The approach is evaluated with a network of Heliomotes over two months during Summer in Los Angeles. To the best of our knowledge, this is the only approach that presents systematic guidelines for offline capacity planning. It will be discussed in more detail in *Sec. 6.3* and used as a baseline for evaluation.

Design experiences of the HydroSolar micro-solar power subsystem are presented in [7]. Despite leveraging the same astronomical model [11] used in this paper, the authors compute the panel size under the assumption that at most 30 minutes of daily charging would have to meet the daily energy requirements, essentially provisioning the solar panel for the worst case scenario (see *Sec. 5*). The battery is selected such that it can support 30 days of operation without harvesting opportunities. However, their deployment did not achieve uninterrupted operation despite seemingly overprovisioning the power subsystem. Interestingly, the authors mention that capacity planning for long-term operation is reasonable, if not necessary, and may be beneficial for improved system utility; however, they do not further investigate this insight.

Another closely related work that relies on the same astronomical model is presented in [22]. The authors empirically validate the model and show that it is, despite its low complexity, very applicable to real-world scenarios. However, the authors are not concerned with capacity planning, but rather focus on runtime prediction of harvesting opportunities with an extended astronomical model. The same battery and panel sizing guidelines as in [7] are used, which may result in an underprovisioned power subsystem.

3 Power Subsystem Design For Long-Term Operation

The ability to achieve uninterrupted long-term operation for solar energy harvesting systems depends on a properly dimensioned power subsystem that can support the expected load. *Fig. 1* shows the high-level design flow of the proposed capacity planning approach for enabling long-term, uninterrupted system operation. The dashed boxes represent user inputs that characterize the deployment setup, hardware technology employed, electrical load, and expected meteorological conditions. The capacity planning




Figure 1: Design flow for power subsystem capacity planning. Dashed boxes and arrows represent user inputs based on which the supported duty-cycle and required battery capacity are computed.

algorithm discussed in *Sec. 5* then computes the battery capacity required to support a user-specified performance level DC_{sys} .

For estimating the theoretically harvestable energy, we leverage the fact that the energy source, the sun, follows both a diurnal and annual cycle. These cycles and the resulting solar energy can be approximated very well with an astronomical model [11]. This model requires deployment location and solar panel setup information as input, i.e. latitude of the deployment site, and orientation and inclination angles of the solar panel, to which we collectively refer as *deployment parameters*.

The astronomical model further depends on three parameters that account for the atmosphere’s optical characteristics and the reflective properties of the ground. However, the exact values for these parameters are highly dependent on time-varying meteorological phenomena. In *Sec. 4.1* we explain how the model parameterization is reduced such that it takes a single input parameter to account for atmospheric and reflective properties. We call this parameter the *environmental parameter* Ω .

The energy that can effectively be harvested on a given day further depends on the *technology parameters*. These parameters characterize the technologies employed by specifying the solar panel’s surface area A_{pv} , its conversion efficiency η_{pv} , and maximum power rating P_{pv} , the efficiencies of the power conditioning circuitry (η_{vr} and η_{cc}), and the charge and discharge efficiencies ($\eta_{bat_{in}}$ and $\eta_{bat_{out}}$) of the chosen storage element. In this work we are not concerned with the selection of optimal storage technology, as this is highly application specific. Note that we leave the selection of panel size to the designer and compute the battery capacity for the given panel. This is because solar panels are available in discrete sizes and scale linearly with output power rating. The capacity of a battery is somewhat less dependent on size. More importantly, however, the panel’s maximum physical size is primarily limited by the form factor of the mote’s enclosure and therefore considered a more sensitive design constraint.

Finally, the *system parameters*, i.e. power dissipation P_{sys} , and minimum expected duty-cycle DC_{sys} , and power conditioning circuitry, i.e., voltage regulator P_{vr} , and charge controller P_{cc} , characterize the load imposed on the battery. The system’s total energy requirement defines the performance level expected by the designer, while the supported duty-cycle $DC_{supported}$, computed by the capacity planning algorithm defines the fraction of the expected performance level that can be sustained with the computed battery capacity B_{nom} . Note that this work assumes constant power dissipation by the load; however, the model also applies to variable power profiles if the average load behavior can be approximated at design time.

In summary, the input parameters discussed in this section characterize the system and expected meteorological conditions such that the energy availability model in *Sec. 4* can accurately approximate the long-term energy harvesting opportunities.

4 System Model

A crucial step in capacity planning consists of estimating the theoretically harvestable energy at a specific point in space and time. To achieve this, we leverage an astronomical energy model [11]. *Sec. 4.1* discusses three modifications to this model such that varying environmental conditions can be taken into account. *Sec. 4.2* describes the harvesting conditioned model, which incorporates conversion and storage inefficiencies.

4.1 Energy Availability Model

According to [11], the total solar energy $E_{astro}(\cdot)$, incident on a flat surface located at latitude L , oriented at azimuth and inclination angles ϕ_p , and θ_p respectively. As shown in (1), it is defined as the sum of the energy contained in direct solar radiation $E_{sun}(\cdot)$, the diffuse radiation by the sky $E_{sky}(\cdot)$, and the reflection of direct and diffuse radiation by the ground $E_{gnd}(\cdot)$, on a given day d , and time of day t [11]. The magnitude of $E_{astro}(\cdot)$, given in $Wh \cdot m^{-2}$, further depends on the distribution and optical characteristics of absorbent gases in the atmosphere, represented by diffuse sky radiation parameter k , optical thickness of the atmosphere τ , and the reflective properties of the ground R , all of which are unit-less. More details on this model and the impact of varying optical characteristics are given in [11], and [23, 24] respectively.

$$E_{astro}(\cdot) = E_{sun}(d, t, L, \phi_p, \theta_p, \tau) + E_{sky}(d, t, L, k, \theta_p, \tau) + E_{gnd}(d, t, L, k, R, \theta_p, \tau) \quad (1)$$

Local obstructions, such as trees, buildings, and meteorological factors (i.e. clouds, snow) also affect the solar energy incident on the panel. Accounting for these effects relies on extensive knowledge of the topographical and meteorological conditions at the deployment site, hence they are not directly considered by the astronomical model. As discussed later in this section, we account for these effects in the calculation of the environmental parameter Ω .

The astronomical model is further expressed in terms of k and τ , both of which are dependent on time-varying optical characteristics of the atmosphere that are difficult to predict [24]. The authors in [11] suggest values of $k = 0$ and $k = 1$ for absolute lower and upper bounds to obtain the contribution by diffuse sky radiation. For the atmosphere's optical thickness, τ , values between 0.1 and 0.4 are recommended, where the former represents a very clear sky, and the latter a very hazy sky [11, 23]. However, for energy harvesting purposes, an upper bound of $\tau = 1$, i.e., no solar harvesting is possible, can be assumed.

In an effort to quantify the parameters k and τ , we note that solar power traces (see *Sec. 6.1*) can be closely approximated with $E_{astro}(\cdot)$ by letting $k = 0.1$ and varying τ , such that $\sum^T E_{astro}(\cdot) \cong \sum^T E_{actual}(d)$. We leverage this observation and define the so called environmental parameter Ω in (2) to replace τ . The parameter δ_i represents the proportion of time during which the atmosphere exhibits the optical thickness τ_i .

$$\Omega = \sum_i^N \delta_i \tau_i, \quad \text{where} \quad \sum_i^N \delta_i = 1 \quad (2)$$

To obtain a representative indication of the atmosphere’s long-term average optical property, the granularity of the weather conditions is represented by N . For example, for a particular geographical location, and with $N = 4$, we might let $\delta = [0.25, 0.35, 0.2, 0.2]$ to represent 25% of the time with clear sky conditions ($\tau_1 = 0.1$), 35% and 20% with light ($\tau_2 = 0.4$), and heavy ($\tau_3 = 0.7$) occlusions respectively, and 20% with no harvesting opportunities at all ($\tau_4 = 1$) (see also *Sec. 6.2* and *6.4*).

As is evident from (1), τ appears in the expressions for E_{sun} , E_{sky} , and E_{gnd} . However, due to fixing k and varying τ , Ω may take values larger than the quantities recommended for τ in the original model. Hence, to minimize the error due to diffuse sky radiation [25], we replace τ with $\min(\Omega, 0.4)$ in the expression for E_{sky} .

The astronomical model assumes a flat horizontal terrain, which, depending on the topography of the deployment site, may not be a valid assumption. The authors in [11] state that the magnitude of $E_{gnd}(\cdot)$ is subject to large error because of topographical variations. However, with a solar panel located at $40^\circ N$, oriented due south (*i.e.*, $\phi_p = 180^\circ$) with 40° inclination angle, and $\tau = 0.2$, $k = 0.3$, the total annual solar energy incident with $R = 0$ is only 2.94% lower than assuming ground reflectivity of bare ground, *i.e.*, $R = 0.3$ [11]. Therefore, unless the effects of ground reflection at a particular deployment site can be obtained through profiling or detailed surface models, it is reasonable to ignore the effect of ground reflection, and assume $R = 0$.

4.2 Harvesting Conditioned Energy Model

The model introduced in the previous section is used to compute the energy incident on a flat surface with a surface area of $1m^2$ for a given time of the year. However, when concerned with electrical energy as opposed to solar energy, various losses due to conversion inefficiencies and self-consumption must be considered [5]. This section discusses the effects of non-ideal harvesting and storage elements.

System Architecture. In this work, a harvest-store-use system architecture as defined in [4] is assumed. In such an architecture, the energy to operate the load is always supplied by the battery, and no bypass path exists that allows operating the load directly from the solar panel when the battery is full, and surplus energy is available. Since there is no dependence on the type of energy store employed, alternative harvesting architectures may also be used. We further assume a stationary solar harvesting installation without sun-tracking capabilities, *i.e.* orientation and inclination angles are fixed.

Panel Characteristics. Only a fraction of the solar energy incident on a solar panel is converted to electrical energy. Depending on technology, the conversion efficiency η_{pv} achieves a few percent for thin-film technologies, and exceeds 40% for high-end multi-junction cells [26]. Furthermore, a solar panel has a manufacturer specified maximum output P_{pv} . This is typically given for Standard Test Conditions (STC), hence the peak power output is not an optimal indication of maximum power. Nevertheless, we assume the maximum possible harvested energy over a time period δt to be limited by $\delta t \cdot P_{pv}$.

In the context of WSN application scenarios, it is desirable to keep the solar panel small in size so to match the mote’s housing and meet low cost expectations [5, 6, 27]. Large-scale photo-voltaic installations are usually only used for WSN base stations and experimental units when mains power is not available. Since the energy model is defined in units of energy per square meter, we account for different solar panel sizes by scaling

the total daily electrical energy output by the panel's surface area A_{pv} . Hence, the total harvested energy on a given day d is approximated with (3), where η_{cc} represent the charge controller's efficiency.

$$E_{pv}(d, \Omega) = A_{pv} \cdot \eta_{cc} \cdot \eta_{pv} \sum_{t=0}^{t=23} \min(1hr \cdot P_{pv}, E_{astro}(d, t, L, k, R, \theta_p, \phi_p, \Omega)) \quad (3)$$

Charge Controller Characteristics. Before the energy transformed by the solar panel can be stored in the battery, a fraction $1 - \eta_{cc}$ of the total energy is lost due to the conversion inefficiency imposed by the charge controller, hence the multiplicative factor η_{cc} in (3). Depending on the chosen technology, the conversion efficiency η_{cc} , can range from 50% for low cost controllers, up to 95% for high-end, *i.e.*, Maximum Power-Point Tracking controllers [7]. The choice of technology is very application specific, and there are arguments advocating advanced charge controllers [28, 29], while others argue that, for micro-solar energy harvesting systems, the gain is dwarfed by the energy expenditure of the controller [7]. It is also possible to operate without a charge controller [15], but the lack of over-voltage protection may significantly reduce battery life [29].

Charge controllers often implement a battery protection mechanism, known as low-power load disconnect [30]. If the battery is fully depleted at any point in time, the load will only be re-connected after the battery state-of-charge has reached a certain percentage of B_{nom} . Prolonged downtime due to protection against deep discharge cycles can incur significant performance penalties. Furthermore, deep cycles severely affect the battery health and its expected lifetime, and should therefore be avoided.

Battery Characteristics. The purpose of the battery is to store harvested energy for supporting the electrical load during periods when harvesting is not possible. However, a battery is not a perfect energy storage element. It suffers from a variety of deficiencies that depend on the battery's chemistry, temperature, discharge rate, and fill-level [30].

Hence, to account for charging and discharging inefficiencies of the battery during simulation, the energy flowing into and out of the battery is scaled by the respective efficiency factors $\eta_{bat_{in}}$ and $\eta_{bat_{out}}$, as shown in (4) and (5). The loss in the charging process due to battery internal resistance and electrochemical processes is represented by $\eta_{bat_{in}}$. The factor $1/\eta_{bat_{out}}$ accounts for the fact that only a fraction of the charge transferred into the battery during charging can be recovered when discharging the battery [30].

$$E_{bat_{in}}(d) = \min(\eta_{bat_{in}} \cdot E_{pv}(d, \Omega), \min(0, \eta_{bat_{out}} \cdot B_{nom} - B(d-1) - E_{bat_{out}}(d))) \quad (4)$$

$$E_{bat_{out}}(d) = E_{load}(d)/\eta_{bat_{out}} + E_{leak} \quad (5)$$

Since a battery has a finite capacity, not all energy generated by the panel may actually flow into the battery, as indicated by the $\min(\cdot)$ function in (4). Similarly, to support the energy consumption $E_{load}(d)$, which represents all energy consumers regardless of function (see (7)), the current battery level and harvested energy must exceed the amount to be withdrawn from the battery, *i.e.*, $B(d-1) + E_{bat_{in}}(d) > E_{bat_{out}}(d)$. The battery specific leakage E_{leak} is assumed to be constant [12]. Then, with the battery characterized by (4) and (5), the battery state-of-charge at the end of a given day is obtained with (6). B_{nom} is the manufacturer rated nominal capacity

converted to Watt-hours. For consistency, $B(d) \geq 0 \forall d$, which means that the expected load may not always be sustained. To circumvent this, the designer may overprovision the battery to enable minimal operation, *i.e.*, $B(d) \geq B_{min} \forall d$.

$$B(d) = \min(\eta_{bat_{out}} \cdot B_{nom}, B(d-1) + E_{bat_{in}}(d) - E_{bat_{out}}(d)) \quad (6)$$

We ignore aging effects of the battery [30]. However, we note that our approach results in very shallow discharge cycles and so protects the expected battery lifetime [30]. In fact, our method results in one full discharge cycle per year; assuming a battery rated for a few hundred discharge cycles [31], the battery will clearly outlast the electronics.

Load Model. The electrical load on the battery consists of all energy consumers present in the system. In addition to the electronic system that performs a particular task, the consumers may include power conditioning, and other supervisory circuitry. For the purpose of capacity planning, the load is specified as the system's total average power dissipation that must be supported by the battery. It is obtained by summing the products of the M system components' duty-cycle (DC_i) and power dissipation (P_{sys_i}). The total daily energy required to operate at the expected performance level is then defined by (7), where $\gamma = 24$ hours.

$$E_{load}(d) = \gamma \cdot \left[P_{cc} + P_{vr} + \sum_i^M (DC_{sys_i} \cdot P_{sys_i}) \right] \quad (7)$$

The power dissipation by the charge controller and input voltage regulator is represented by P_{cc} and P_{vr} , respectively. These are assumed to be always operational. Depending on design optimizations, however, both of them may be duty-cycled to reduce energy consumption. In that case, their average power dissipation is computed identically to that of the system components.

5 Capacity Planning for Long-Term Uninterrupted Operation

As discussed in *Sec. 4*, the harvestable energy can be closely approximated if setup and technology parameters are known. To account for the effects of meteorological conditions, and thus more closely approximate the long-term energy input, the environmental parameter Ω was defined. This parameter can be obtained with (2), or, if available, by profiling a representative dataset. The total energy consumption, defined in (5), completes the necessary information for long-term capacity planning.

Intuitively, a battery should be sized exactly such that (i) it can support the expected operation during periods of solar energy deficit, and (ii) be replenished by the panel during times of solar energy surplus. We consider the annual solar cycle to compute the power subsystem capacity such that uninterrupted long-term system operation with a minimum battery capacity and solar panel size can be achieved. Two such cycles are illustrated in *Fig. 2*, which shows the actual energy input $E_{actual}(d)$ of the CA dataset, the daily system consumption to be supported, and the model approximation $E_{bat_{in}}(d)$ such that $\sum^T E_{actual}(d) \cong \sum^T E_{bat_{in}}(d)$ over $T = 720$ days.

The battery capacity, B_{nom} , required to operate the system during times of deficit, *i.e.*, $E_{bat_{out}}(d) > E_{bat_{in}}(d)$, $d \in [d_1, \dots, d_2]$, is given in (8) and illustrated by the cross-hatched area in *Fig. 2*. Note that we assume the battery to be fully charged on day d_1 . The first term on the left-hand side in (8) specifies the amount of energy that is necessary




Figure 2: Actual (E_{actual}), and modeled daily energy input ($E_{bat_{in}}$) and output ($E_{bat_{out}}$) for panel area $A_{pv} = 0.02m^2$ over two years of the CA dataset (see Sec. 6.1). Surplus energy is indicated with the hatched area. The cross-hatched area shows the required battery capacity.

to support the system operation, while the second term represents the modeled energy input expectations. The difference is then the minimum required battery capacity B_{nom} .

$$\sum_{d_1}^{d_2} (E_{bat_{out}}(d) - E_{bat_{in}}(d)) \leq B_{nom} \quad (8)$$

Similarly, (9) specifies the amount of energy that is harvested in excess of what is required to sustain short term operation during periods of surplus, *i.e.*, over the interval $[d_0, \dots, d_1]$ (hatched area in Fig. 2). As mentioned previously, the harvested energy during periods of surplus must be able to recharge the battery. To achieve perpetual operation over multiple years, the left-hand term in (9) must therefore be at least as large as the left-hand term in (8).

$$\sum_{d_0}^{d_1} (E_{bat_{in}}(d) - E_{bat_{out}}(d)) \geq B_{nom} \quad (9)$$

The required battery capacity B_{nom} can then be obtained by finding the intersections d_0 , d_1 , and d_2 between $E_{bat_{in}}(d)$ and $E_{bat_{out}}(d)$ such that inequalities (8) and (9) hold. For example, with the CA dataset shown in Fig. 2, and a panel size of $20cm^2$ and $E_{load} = 227.8mW$, the required battery capacity is approximately 68Wh (see Sec. 6.3). When reducing the panel size by 50%, the battery capacity must be increased by roughly 67% in order to achieve a performance level of 65% of the larger panel setup. This clearly shows the non-linear relationship between battery capacity and solar panel performance.

6 Evaluation

This section evaluates the proposed method's ability to yield a battery capacity that ensures uninterrupted long-term operation of solar energy harvesting systems.

6.1 Evaluation Methodology, Validation Data, and Performance Metrics

Methodology. To validate the modified astronomical model, we first compare its energy estimations to measurements performed with a solar panel. As a second step, we evaluate the model's ability to support long-term operation by simulating a system according to Sec. 4 and 5, and the simulation input data discussed in the following.

Model Validation Input Data. For validation of the modified astronomical model, we obtain ground-truth data by measuring the power generated by a $0.1725m^2$ monocrystalline solar panel (cleversolar CS-30) rated at 30 Watt over a period of 41 days (22/07/2013 - 08/31/2013). The power generated by the panel, and dissipated over a purely resistive load was sampled at $1Hz$ with a custom measurement circuit. The panel

Table 1: Name, time-period, and location of NSRD¹ datasets used for evaluation of the proposed approach. Maximum, mean, minimum and variance of solar radiation are given in $Wh/0.01m^2$.

Name	Time Period	Latitude	Longitude	Maximum	Mean	Minimum	Variance
CA	01/01/99 – 12/31/09	34.05	-117.95	10.37	7.03	0.92	5.62
MI	01/01/99 – 12/31/09	42.05	-86.05	10.55	5.34	0.53	9.05
ON	01/01/99 – 12/31/09	48.05	-87.65	10.98	5.07	0.44	11.24

was placed on the roof-top of our university building at $47.37^\circ N$, $8.55^\circ E$, and oriented with azimuth, and inclination angles of 170° and 70° respectively. This particular location has clear view of the sky without any obstructions that could lead to shading, hence deviations from the model can be assumed to originate from weather effects only. Ground reflections are assumed to be negligible.

Simulation Input Data. For the simulation input data, we resort to the National Solar Radiation Database¹ (NSRD) from where we obtain hourly, global (i.e. direct and diffuse) solar radiation for three locations in California (CA), Michigan (MI), and Ontario (ON) (see *Table 1*). We use 11 years of data, from which the first year (i.e., days 1-365) of each location is used as calibration data (see *Sec. 6.3*), while the data for the remaining 10 years is used as input for the simulation discussed in *Sec. 6.3*.

The data traces from NSRD are given in $Wh \cdot m^{-2}$ of solar energy incident on a flat surface with zero inclination. To account for smaller panel sizes, inefficiencies of individual components, and losses in energy storage during simulation, the data is conditioned as explained in *Sec. 4.2*. For the technology parameters we assume $\eta_{bat_{in}} = 0.9$, $\eta_{bat_{out}} = 0.7$, $\eta_{cc} = \eta_{vr} = 1$, $\eta_{pv} = 10\%$, and zero inclination angle. The reconnect hysteresis (see *Sec. 4.2*) is set at 30% of battery capacity $B = \eta_{bat_{out}} \cdot B_{nom}$.

Performance Metrics. For performance comparison between the approach discussed in this paper and the state-of-the-art (*SotA*) [12], we define the following metrics. Since size and cost considerations play a major role in *WSN* scenarios, an optimal energy harvesting system is one with the smallest hardware configuration that is able to achieve the expected performance level, i.e., when E_{load} is sustained over the entire simulation.

For evaluating the performance of the two approaches, we assume that the maximum feasible panel size is given, and wish to obtain the minimum battery capacity such that uninterrupted operation can be achieved. We report the computed battery capacities, and percentage of time spent with depleted battery for each configuration and dataset. Any set of input parameters that cannot support long-term operation at the expected performance level is considered invalid.

6.2 Modified Astronomical Model Validation

The measurement data described in *Sec. 6.1* is used as ground truth for validation of the modified astronomical model. Here we are concerned with how well the actual energy input can be approximated with a given Ω . We assume $\Omega = 0.51$ as per (3) with $\tau = [0.1, 0.4, 0.7, 1]$ and $\delta = [0.25, 0.35, 0.2, 0.2]$ to represent the expected weather condition. For the panel efficiency we use $\eta_{pv} = 21.5\%$ according to the specification, and the technology and deployment parameters from *Sec. 6.1*.

As is evident from *Fig. 3*, which shows the ratio of estimated and actual energy input for each day, the model tends to significantly overestimate available energy for days

¹ http://rredc.nrel.gov/solar/old_data/nsrdb/1991-2010




Figure 3: Ratio of total daily energy predicted with our approach ($E_{pv}(d, \Omega = 0.51)$), and actual energy input ($E_{actual}(d)$) for each day. Note: bars are capped at 2, and labeled with actual value.

Table 2: Modified astronomical model validation statistics.

(a) Statistics for 41 day experiment.		(b) Weekly energy sums. E_{actual} refers to measurements with solar panel.									
Min:	0.7947	Mode:	0.7947								
Max:	26.87	St. Dev.:	5.737								
Mean:	3.389	Range:	26.07								
Median:	1.01										
				Week							
				1	2	3	4	5	6	Total	
				E_{pv} [Wh]	4844	5646	5034	5109	5164	3706	29503
				E_{actual} [Wh]	3773	5775	3155	4929	3605	3255	24492
				E_{pv}/E_{actual}	1.28	0.98	1.6	1.04	1.43	1.14	1.20

with little, to no energy harvesting opportunities. *Table 2a* lists statistics for the same 41-day period. *Table 2b* shows the results when considering an estimation granularity of one week. Over the entire period, the actual conditions are overestimated by 20.04%. This is a good result, considering that even more elaborate models tend to suffer a great deal from uncertainties, particularly due to modeling of diffuse sky radiation [25].

6.3 Capacity Planning Performance Evaluation

The previous section showed that our energy availability model can closely approximate actual conditions. This section now demonstrates that our approach in fact supports uninterrupted long-term operation of a simulated system for a variety of input data. Aside from *SotA* [12], we are not aware of concrete algorithms for power subsystem capacity planning of energy harvesting systems. Thus, the method described in [12, 32, 33] is used as a baseline and briefly reviewed in the following.

Reference Model. In [12] the authors present a set of abstractions for capacity planning of energy harvesting systems that can be considered the state-of-the-art (*SotA*) in harvesting theory. The authors define Energy-Neutral Operation (ENO) as a performance metric, and formally state the conditions that must be met to achieve ENO. The authors argue that a system’s total average power dissipation, ρ_c , must always be less than, or equal to the source’s average power generation, ρ_s . If energy inefficiency is acceptable, *i.e.*, dissipating the power generated by the panel as heat when the battery is full, the minimum battery capacity is defined by the sum of the maximum negative deviation from ρ_s , and the maximum positive deviation from ρ_c . If wasting is not permitted, the battery capacity must be increased by the maximum positive deviation from ρ_s such that surplus energy can be buffered [32]. Despite defining a capacity for long-term continuous operation, the authors conclude that, for achieving ENO, the battery state-of-charge $B(d)$ on day d must be no less than $B(d - 1)$. With this approach the benefits of capacity planning are not fully leveraged; the system must rely on a well performing energy prediction scheme to achieve acceptable long-term performance.

In order to extract ρ_s , the *SotA* algorithm requires a dataset that is representative of the conditions at the deployment site. However, when applying the method described in [32] to obtain this quantity, it is found that their approach yields significantly different performance levels depending on the particular time period used. In fact, even

Table 3: Expected, and actual performance level achieved (in mW), and required battery capacity (i.e., $B = \eta_{bat_{out}} \cdot B_{nom}$, in Wh) obtained from simulation of 4 panel sizes (in m^2) and the 3 datasets listed in Table 1 for *SotA* and the method proposed in this work.

	A_{pv}	MI			CA			ON		
		Expected	Actual	B	Expected	Actual	B	Expected	Actual	B
<i>SotA</i>	0.005	55.73	52.32	42.51	73.17	73.17	61.48	56.75	51.04	56.96
This work			55.73	88.85		73.17	56.84		56.75	137.5
<i>SotA</i>	0.01	111.46	104.65	85.03	146.28	146.28	122.96	113.53	102.11	113.93
This work			111.46	168.39		146.28	113.5		113.53	275.1
<i>SotA</i>	0.015	157.54	148.34	106.44	197.85	197.85	133.89	159.1	145.19	171.42
This work			157.54	218.74		197.85	110.8		159.1	363.81
<i>SotA</i>	0.02	186.71	179.18	113.9	227.83	227.83	127.02	186.72	173.01	180.58
This work			186.71	216		227.83	67.59		186.72	381.6

for datasets with little variance, e.g., the CA dataset (see Table 1), ρ_s converges to its average value only after a few seasonal cycles. This exemplifies that, when attempting to achieve uninterrupted long-term operation at a predefined performance level, consideration of the source’s longest-term cycle is necessary. In the case of solar harvesting, the period can generally be assumed to be one year. However, in areas where significant meteorological phenomena with a periodicity of multiple years occur, improved results may be obtained if the analysis is performed over the respective period.

Simulation. The *SotA* approach attempts to compute the supported performance level and required battery size based on a given power profile. Our approach, on the other hand, takes the expected performance level as input. Hence, to evaluate and compare the two approaches through simulation, we first obtain the respective performance levels and battery capacities as follows. With the setup and technology parameters defined in Sec. 6.1, we find the battery capacity and supported performance level with the *SotA* approach and one year of calibration data for all three datasets listed in Table 1 and four panel sizes, i.e., $5cm^2$, $10cm^2$, $15cm^2$, and $20cm^2$. The respective performance levels obtained with *SotA* are then used as input to our model to compute the minimal capacities required for each dataset. Once these quantities have been found for *SotA* and our approach, we run a simulation with the remaining 10 years of data.

Note that our approach does not require calibration data. However, since *SotA* relies on a representative power trace, we allow our approach to extract the weather conditions from the calibration data to compute Ω (see Sec. 4.1). We use $\tau = [0.1, 0.4, 0.7, 1]$ and $N = 4$ in (2) and let δ_i be the days with more than 75%, 50%, 25%, and 0% of the maximum expected energy, i.e., $E_{astro}(\Omega = 0.1)$. Since meteorological conditions tend to follow a certain periodicity, Ω may be obtained with very little data. In fact, we did not find significant improvements when using more than half a year of calibration data, as long as the data is representative of the conditions during the critical periods of continuous solar energy deficit, i.e., winter in the northern hemisphere.

Results. The results obtained from simulation are summarized in Table 3, and discussed in the following. As is evident, the *SotA* approach achieves an acceptable performance only for the CA dataset. As defined in Sec. 6.1, acceptable performance means that the battery can support the user-specified performance level indefinitely.

While the *SotA* approach yields smaller battery capacities for the MI and ON datasets compared to our approach, the configurations fail to sustain the expected performance

level over the entire simulation period. Our approach, on the other hand, achieves the expected performance level with zero down-time for all configurations and simulations.

CA Dataset. With the CA dataset, *SotA* achieves the expected performance level for the entire 10 year period with all simulated panel sizes. This comes at no surprise; the authors of the *SotA* approach are located in Southern California and used locally measured data for design and verification of their approach. From *Table 1* we see that this particular dataset has the lowest input data variance, with roughly half of that of the other datasets. Nevertheless, with $A_{pv} = 0.005m^2$ and $A_{pv} = 0.01m^2$, the *SotA* approach yields a battery size that is about 1.48 times the size of the minimal capacity. For the other two panel sizes, the algorithm overestimates the absolute minimal possible capacity by a factor of 1.86, and 2.7, respectively.

For the same dataset, our approach yields smaller battery capacities that can sustain the expected performance level over the entire simulation period. When compared to *SotA*, a reduction of roughly 8% in capacities are obtained with $A_{pv} = 0.005m^2$ and $A_{pv} = 0.01m^2$. For $A_{pv} = 0.015m^2$ and $A_{pv} = 0.02m^2$ our approach yields 17.25% and 46.8% smaller capacities than *SotA*. This is an interesting result because it shows that average generation, ρ_s , is not a good indicator of the long-term sustainable performance level. With increasing panel size, ρ_s behaves in a manner that may not be representative of the long-term dynamics, causing the *SotA* model to assume an overly pessimistic negative deviation from ρ_s , and yield a larger capacity than necessary.

MI and ON Datasets. The results in *Table 3* show that *SotA* does not achieve satisfactory performance for the MI and ON dataset. Hence, we only focus on the results of our approach. For the MI dataset, the minimal possible battery capacity is overestimated by a maximum of 13.5%. For the ON dataset, our approach overestimates by up to 22.1%. This constant, but reasonable overestimate is due to assuming Ω larger than absolutely necessary. In fact, reducing Ω by 20% for ON leads to an overestimate of only 1.4%. For the MI dataset, Ω must be reduced by 7% to achieve the same result. However, doing so would result in dangerously low battery fill-levels during times of deficit, causing the system to become susceptible to low-power disconnect penalties.

When given the opportunity to analyze two full years of calibration data, the *SotA* approach fails to achieve the expected performance level only for one of the configurations of the MI dataset. With $A_{pv} = 0.015m^2$, the system spends 0.59%, or 21.5 days of the time with a drained battery. Considering that *SotA*'s performance depends on the closeness of ρ_s to actual average generation, longer data traces are expected to improve its performance. Interestingly, *SotA* overestimates the battery capacities for the CA dataset almost identically when only one year of data is available, which is due to the low variance in the energy input. For the other two datasets, both approaches yield comparable capacities, despite giving *SotA* the advantage of analyzing two years of calibration data, while our method only used one year of calibration data.

Energy Approximation. In this work we focused on long-term provisioning because short-term deviations from the model should be absorbed by a properly sized battery. Here we investigate how our model can cope with source variations that lead to energy deficit. *Fig. 4* shows the ratio of total energy approximated by the model, and effectively harvested energy for each year on the left, and the same for the periods of continuous deficit on the right. As is evident, with the largest panel size used ($A_{pv} = 0.02$), the




Figure 4: Ratio of total energy E_{pv} approximated with our approach, and actual energy input (E_{actual}) for each year (left), and for each year's period of deficit (right). Note the scales.




Figure 5: Achieved performance level for CA (left), MI (middle), and ON (right) datasets and four panel sizes with down-scaled energy input.

model assumes on average around 95% of the actual annual energy input for the CA and MI datasets, and roughly 88% for ON. The ratios of approximated and actual energy input for the panels with area 0.005 m^2 and 0.01 m^2 are about 85% for CA and MI, and around 78% for ON.

For the periods of deficit, shown on the right hand graphs in *Fig. 4*, much more variation is evident from one year to the next. However, the approximations with different panel sizes are much less scattered, and the approximation with the larger panel size is not always best. Nevertheless, on average, the model underestimates actual conditions by roughly 20% for the CA and MI datasets, and about 25% for ON. This shows that the battery obtained is reasonably overprovisioned, and will be able to safely bridge short periods with energy input below modeled long-term expectations.

6.4 Sensitivity Analysis

In this section we investigate the model's sensitivity to the selection of the environmental parameter Ω . We further exemplify the importance of choosing inclination, and azimuth angles such that they are representative of the actual deployment setup.

Environmental Parameter. As is evident from *Sec. 6.2*, the model's ability to accurately estimate the long-term expected energy input depends on proper choice of Ω . This parameter is used to account for environmental effects due to *e.g.*, meteorological conditions and local obstructions. These cause the model to overestimate actual energy input, which is equivalent to assuming too low of a value for Ω .

Hence, to investigate these effects on the model's estimation accuracy, we scale Ω with a scaling factor S (*i.e.*, $E_{pv}(\dots, S \cdot \Omega)$ in (4)), and simulate the system as discussed in *Sec. 6.1*. The achieved mean duty-cycle over ten years for the three datasets are shown in *Fig. 5*. The results show that the expected duty-cycle can be achieved when the energy input is at least 75%, 85%, and 90% of the original magnitude for the CA, ON, and MI datasets respectively.

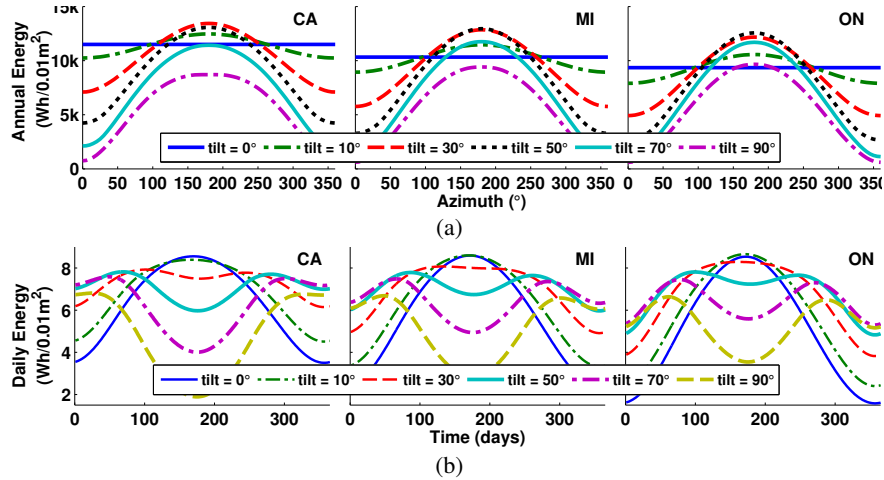


Figure 6: (a) Total *annual* energy incident with inclination angles from 0° to 90° as a function of azimuth angle. (b) Total *daily* energy incident over one year for inclination angles ranging from 0° to 60° ($\tau = 0.1$, $k = 0.3$, $R = 0.3$, $\eta_{pv} = 0.1$).

Setup Parameters. The effects of varying the setup parameters, i.e. inclination and orientation angles, on the harvestable energy are illustrated in *Fig. 6a*. It shows the total *annual* energy incident on a panel with a surface area of 10cm^2 over the course of one year for the three datasets in *Table 1* and various inclination angles as a function of the panel orientation. As is intuitively clear, for an inclination angle of 0° , i.e. the panel is placed parallel to the ground, the orientation has no effect.

Similarly, *Fig. 6b* shows the total *daily* energy incident on the same panel for various inclination angles. The effect on the harvestable energy due to setup parameters, and seasonal variations is clearly visible, solidifying our argument that the source’s seasonal behavior must be considered when attempting to achieve uninterrupted long-term operation at a predefined performance level.

7 Conclusions

In this work we presented a systematic approach to offline capacity planning of the power subsystem for solar energy harvesting systems. The approach is based on a modified astronomical model and takes into account seasonal variations of the energy source to enable uninterrupted long-term operation. Solar power measurements of a real panel are used to validate the modified model. We further compared our approach to the state of the art (*SotA*) in harvesting theory through simulation with real-world input data, and showed that the proposed method achieves zero down-time (compared to up to 10% for *SotA*) for three different locations and four different panel sizes while requiring up to 53% smaller batteries. The results show that pre-deployment design considerations are absolutely inevitable for achieving long-term uninterrupted system operation. In order to enable the system to adapt to significant deviations from the model, and therefore improve the energy efficiency, we are currently extending this work with a low-complexity, power subsystem aware, dynamic duty-cycling scheme.

Acknowledgments. This work was supported with a grant from the Swiss Nano-Tera.ch initiative and evaluated by the Swiss National Science Foundation.

References

1. Mini, R.A., Loureiro, A.A.: Energy in wireless sensor networks. In: *Middleware for Network Eccentric and Mobile Applications*. Springer (2009) 3–24
2. Buchli, B., et al.: GPS-equipped wireless sensor network node for high-accuracy positioning applications. In: *Wireless Sensor Networks*. Springer (2012) 179–195
3. Chalasani, S., Conrad, J.M.: A survey of energy harvesting sources for embedded systems. In: *Southeastcon, 2008*. IEEE, IEEE (2008) 442–447
4. Sudevalayam, S., Kulkarni, P.: Energy harvesting sensor nodes: Survey and implications. *Communications Surveys & Tutorials, IEEE* **13**(3) (2011) 443–461
5. Hanssen, L., Gakkestad, J.: Solar Cell Size Requirement for Powering of Wireless Sensor Network Used in Northern Europe. In: *Proceedings of the International Workshops on PowerMEMS*. (2010) 17–20
6. Seah, W.K., et al.: *Research in Energy Harvesting Wireless Sensor Networks and the Challenges Ahead*. (2012)
7. Taneja, J., et al.: Design, Modeling, and Capacity Planning for Micro-solar Power Sensor Networks. In: *Proceedings of the 7th international conference on Information processing in sensor networks*. IPSN '08, Washington, DC, USA, IEEE Computer Society (2008) 407–418
8. Le, T.N., et al.: Power Manager with PID controller in Energy Harvesting Wireless Sensor Networks. In: *Green Computing and Communications (GreenCom), 2012 IEEE International Conference on*, IEEE (2012) 668–670
9. Vigorito, C.M., 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*, IEEE (2007) 21–30
10. Piorno, J.R., 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*, IEEE (2009) 6–10
11. Dave, J., et al.: Computation of Incident Solar Energy. *IBM Journal of Research and Development* **19**(6) (1975) 539–549
12. Kansal, A., et al.: Power management in energy harvesting sensor networks. *ACM Transactions on Embedded Computing Systems (TECS)* **6**(4) (2007) 32
13. Zhang, P., et al.: Hardware design experiences in ZebraNet. In: *Proceedings of the 2nd international conference on Embedded networked sensor systems*. SenSys '04, New York, NY, USA, ACM (2004) 227–238
14. Raghunathan, V., et al.: Design considerations for solar energy harvesting wireless embedded systems. In: *Proceedings of the 4th international symposium on Information processing in sensor networks*, IEEE Press (2005) 64
15. Park, C., Chou, P.H.: Ambimax: Autonomous energy harvesting platform for multi-supply wireless sensor nodes. In: *Sensor and Ad Hoc Communications and Networks, 2006. SECON'06. 2006 3rd Annual IEEE Communications Society on*. Volume 1., IEEE (2006) 168–177
16. Sitka, P., et al.: Fleck-a platform for real-world outdoor sensor networks. In: *Intelligent Sensors, Sensor Networks and Information, 2007. ISSNIP 2007. 3rd International Conference on*, IEEE (2007) 709–714
17. Glatz, P.M., et al.: Designing perpetual energy harvesting systems explained with rivermote: A wireless sensor network platform for river monitoring. *Electronic Journal of Structural Engineering, Special Issue: Wireless Sensor Networks and Practical Applications* (2010) 55–65
18. Lu, J., Whitehouse, K.: 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*, ACM (2012) 245–256
19. Sharma, N., et al.: Predicting solar generation from weather forecasts using machine learning. In: *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on*, IEEE (2011) 528–533
20. Bacher, P., et al.: Online short-term solar power forecasting. *Solar Energy* **83**(10) (2009) 1772–1783
21. Kooti, H., et al.: Energy Budget Management for Energy Harvesting Embedded Systems. In: *Embedded and Real-Time Computing Systems and Applications (RTCSA), 2012 IEEE 18th International Conference on*, IEEE (2012) 320–329
22. Jeong, J., Culler, D.: Predicting the Long-Term Behavior of a Micro-Solar Power System. *ACM Transactions on Embedded Computing Systems (TECS)* **11**(2) (2012) 35
23. Bohren, C.F., Clothiaux, E.: *Atmospheric Optics*. Wiley-VCH (2006)
24. Heinemann, D., 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)
25. Gubler, S., Gruber, S., Purves, R.: Uncertainties of parameterized surface downward clear-sky shortwave and all-sky longwave radiation. *Atmospheric Chemistry and Physics* **12**(11) (2012) 5077–5098
26. Green, M.A., Emery, K., Hishikawa, Y., Warta, W., Dunlop, E.D.: Solar cell efficiency tables (version 39). *Progress in photovoltaics: research and applications* **20**(1) (2011) 12–20
27. Pister, K. In: *Smartdust: Autonomous sensing and communication in a cubic millimeter*
28. Brunelli, D., et al.: An efficient solar energy harvester for wireless sensor nodes. In: *Proceedings of the conference on Design, automation and test in Europe*, ACM (2008) 104–109
29. Corke, P., et al.: Long-duration solar-powered wireless sensor networks. In: *Proceedings of the 4th workshop on Embedded networked sensors*. EmNets '07, New York, NY, USA, ACM (2007) 33–37
30. Buchli, B., et al.: Battery state-of-charge approximation for energy harvesting embedded systems. In: *Wireless Sensor Networks*. Springer (2013) 179–196
31. Bergveld, H.J.: *Battery management systems : design by modelling*. PhD thesis, Enschede (June 2001)
32. Kansal, A., Potter, D., Srivastava, M.B.: Performance aware tasking for environmentally powered sensor networks. In: *ACM SIGMETRICS Performance Evaluation Review*. Volume 32., ACM (2004) 223–234
33. Kansal, A., et al.: Harvesting aware power management for sensor networks. In: *Proceedings of the 43rd annual Design Automation Conference*, ACM (2006) 651–656