

# Policies for Predictive Energy Management with Supercapacitors

Christian Renner, Florian Meier, Volker Turau  
*Institute of Telematics*  
*Hamburg University of Technology, Hamburg, Germany*  
*Email: {christian.renner,turau}@tu-harburg.de*

**Abstract**—This paper presents an algorithm to dynamically determine the maximum supported uniform demand for energy of sensor nodes powered by energy harvesters using supercapacitors as energy buffers. Knowledge about the maximum uniform consumption is required to adapt the sensor node's duty cycle or task schedule to achieve uniform, utility-maximizing, and depletion-safe operation. Our algorithm makes use of a supercapacitors' relationship between state-of-charge and voltage, is particularly designed to handle the non-linear system model, and is lightweight enough to run on low-power sensor node hardware. We define three energy policies, evaluate their performance using a real-world solar-harvesting trace, and analyze the influence of the supercapacitor's capacity and errors of the energy forecast.

## I. INTRODUCTION

Energy harvesting offers perpetual and maintenance-free operation for sensor nodes [1], [2]. It is yet difficult to align a node's power consumption (i.e., its task schedule or radio duty cycle) with the energy produced by a harvester a priori, so as to achieve depletion-safe while uniform and utility-maximal operation. Many harvesters, e.g., solar cells, show an unsteady production profile: Energy intake varies within orders of magnitude, may abruptly change, and depends considerably on a node's location. The maximum supported uniform consumption of a node can thus be hardly foreseen.

Opportunistic and greedy consumption (i.e., take what you get) is not an option. Many applications require uniform and reliable node operation to guarantee constant sampling rates and achieve low latency and connectivity in multi-hop environments. Several solutions to this problem have been presented [3]–[5]. They either require a linear system model or do not work well for low-capacity energy buffers. Thus, they cannot be employed on many low-cost energy-harvesting sensor nodes that solely use supercapacitors [2], [6]. These energy buffers yet have numerous advantages over batteries, among these being simple (and low-cost) charging circuits, virtually unlimited charging cycles, and the direct relationship between state-of-charge and voltage [6], [7].

In this paper, we present and evaluate the concept of energy policies to enable load-adaptation for this kind of platform. The idea behind this concept is to simulate the node's future energetic course using an energy-harvest forecast and to identify the maximum load complying with the energy policy. The maximum load can be used to adapt a nodes radio-duty cycle or can be utilized as an energy

budget [8]. An energy policy defines state-of-charge requirements within the prediction horizon. Since the definition of energy policies is tightly coupled with the harvester hardware, we introduce our hardware platform. We derive a system model for its energy flow and an algorithm to simulate the node's energetic future for a given energy-harvest forecast. We present three energy policies and show a light-weight algorithm for load-adaptation. We compare the policies using a real-world energy-harvest trace and analyze the influence of the energy buffer's capacity and energy-harvest forecast errors. We discuss limitations of our approach and address the remedies for the weaknesses of our policy-based approach and give pointers to future work.

## II. RELATED WORK

Several energy-harvesting power supplies for sensor nodes exist. Prometheus [1] is based on a two-stage energy storage system consisting of a supercap as the primary energy source and a rechargeable Li<sup>+</sup> battery. A solar cell serves as energy source. Charging and discharging behavior of the circuit and supercap are examined, while taking a first step into the direction of energy-aware scheduling by adapting the duty cycle to the current supercap voltage. Prometheus has been successfully deployed in [9]. The Everlast platform stores energy harvested by a solar cell in a supercap solely [2]. Maximum power point tracking (MPPT) is employed to increase the produced electrical power. The harvester used in this paper is a cross section of all presented works.

Kansal et al. presented a definition of energy-neutral operation of sensor networks [4]. They have developed abstractions to characterize the complex time varying nature of self-sustaining energy sources with analytically tractable models including an energy prediction model. Moser et al. [3] have presented tools and methods that use multiparametric programming achieving optimal performance in energy harvesting systems. Simulational results and the evaluation of the algorithms on real hardware are discussed. The LQ-tracker presented in [5] tries to adapt a node's duty cycle to its current battery level. This concept makes no assumptions about the harvesting source, but does not function well for supercaps, which offer a rather low capacity and operate over a wide voltage range.

Algorithms for the generation of energy-intake forecasts have been studied. EWMA filtering with time slots exploits

the cyclic pattern of harvesters, such as solar cells [4]. The Weather-Conditioned Moving Average (WCMA) [10] to combining the history of the last days and the current weather situation to improve short-term forecasts.

### III. ENERGY-HARVESTING SENSOR NODE

The energy-flow model and energy policies presented in this paper are based on a customized energy-harvesting power supply, the *harvester*, for the Iris node.

A solar cell with maximum power-point at 4 V and 35 mA serves as a harvester. We chose a direct charging circuit instead of a maximum power-point tracker as in [2]. The advantages are a simple, low-cost circuit and a charging current only depending on the lighting conditions. The main disadvantage is that the harvested energy depends on the supercap voltage. Our hardware platform provides a sensor for measuring the current produced by the solar cell.

The harvester is designed for supercaps with a maximum voltage of  $V_{\max} = 2.7$  V and, generally, any available capacity. Results from the literature [1], [6] show that supercaps with 25 to 100F offer a small size and cheap prize. They operate a sensor node for a few days at a 1% radio duty cycle. The node can access the supercap's voltage and thus state-of-charge via an ADC port. To protect the supercap from overcharging, the harvester automatically disconnects the supercap from the solar cell, if its voltage exceeds  $V_{\max}$ .

We chose the Texas Instruments TPS 61220 to supply the sensor node with a constant voltage of  $V_n = 2.7$  V. The regulator starts operation when  $V_c$  exceeds 1.6 V for the first time and has a cut-off voltage of  $V_{\text{cut}} = 0.5$  V on our target platform. We measured a switching efficiency  $\eta$  of 75% to 95% depending on input voltage and output load.

### IV. SYSTEM MODEL: SIMULATION AND PREDICTION

From the simplified harvester circuit displayed in Fig. 1, we derive a mathematical model encompassing the three building blocks (energy reserve, intake, consumption):

$$\eta \cdot V_c \cdot (I_h - I_c) = V_n \cdot I_n. \quad (1)$$

Modeling the relationship of supercap voltage  $V_c$  and current  $I_c$  as an ideal capacitor, i.e.,  $I_c = C \cdot \dot{V}_c$ , we obtain

$$C \cdot \dot{V}_c = I_h - \frac{I_n \cdot V_n}{\eta \cdot V_c}. \quad (2)$$

We neither consider leakage nor charging effects, as they are negligible w.r.t. the consumption of sensor nodes [6], [7].

#### A. Simulating a Node's Energetic Course

It is possible to simulate a node's energetic course in terms of  $V_c$  using (2) for a given current consumption  $I_n$  and harvester intake  $I_h$  but requires simplifications. Solving (2) analytically is impossible, and the hardware characteristics imply additional constraints: The overcharging-protection prevents  $V_c$  from exceeding  $V_{\max}$ , and the switching regulator fails if  $V_c < V_{\text{cut}}$  (cf. Sect. III).

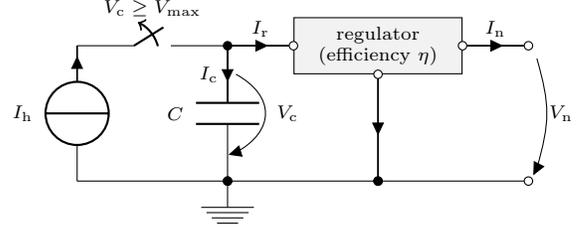


Figure 1. Simplified equivalent circuit of the harvester

A fine-grained numerical evaluation is possible yet computationally infeasible, since the precise temporal course of  $I_n$  and  $I_h$  is required. Due to these limitations, we simplify the problem by assuming that  $I_h$  and  $I_n$  are piecewise constant functions. This simplification complies with reality:

- Only discrete samples of a real harvester intake trace  $I_h$  exist. These samples can be converted into a piecewise constant function of  $I_h$ , but do not yield the precise course of  $I_h$  at any possible time.
- The (average) consumption of a node changes infrequently and is constant in small time intervals.

For  $I_h = \text{const.}$ ,  $\eta = \text{const.}$ ,  $I_n = \text{const.}$ , (2) reduces to the ordinary, first-order differential equation (ODE)

$$\dot{y} = b - \frac{a}{y} \quad \left( y = V_c, a = \frac{V_n \cdot I_n}{\eta \cdot C} \geq 0, b = \frac{I_h}{C} \geq 0 \right). \quad (3)$$

This equation has the implicit solution

$$0 = y - y_0 + \frac{a}{b} \cdot \log \left( \frac{a - b \cdot y}{a - b \cdot y_0} \right) - b \cdot \Delta t = f(y, y_0, \Delta t), \quad (4)$$

where  $y = y(t)$  and  $y_0 = y(t_0)$  with  $t \geq t_0$  being points in time with  $\Delta t = t - t_0$

This equation can be solved with, e.g., Newton's Method

$$y_{n+1} = y_n - \frac{f(y_n, y_0, \Delta t)}{f'(y_n, y_0, \Delta t)} = g(y_n, y_0, \Delta t) \quad (5)$$

for a given number of iteration steps or until an absolute or relative error  $\epsilon$  is achieved. Due to the overcharging protection,  $y = V_c \leq V_{\max}$  must be enforced during calculation. Moreover,  $I_n = 0$  if  $V_c < V_{\text{cut}}$  (cf. Sect. III).

In two special cases, (4) cannot be used:

- 1) If  $I_h = 0$  and thus  $b = 0$ , we find  $y = \sqrt{y_0^2 - 2a \cdot \Delta t}$
- 2) For  $y_0 = \frac{a}{b}$ , the solution is  $\dot{y} = 0 \Rightarrow y = y_0$

Figure 2 shows the corresponding algorithm.

#### B. Online Prediction of the Energetic Future

The approach used for simulation can be turned into an algorithm for predicting a node's energetic future. For this purpose, an estimate of the harvester intake  $I_h$  is required. State-of-the-art energy intake prediction [4], [11] generates such a forecast in terms of a set of time slots and corresponding mean intake values. Intake  $I_h$  does hence not change within a slot. Assuming a constant average load  $I_n$  in each slot reduces the overhead of predicting the energetic

future by solving (4) with (5), because only one solution has to be calculated for each end of a time slot, and updates have to be calculated only if a slot elapses. The results of [11] suggest that 12 to 24 slots per day, requiring recalculations only once every one or two hours, are sufficient.

## V. POLICY-BASED LOAD MAXIMIZATION

In this section, we present the concept of energy policies and introduce an algorithm to identify the maximum constant load satisfying a given policy. The maximum load can, e.g., be used to set up the node's duty cycle or task schedule.

### A. Energy Policies

An energy policy  $\mathcal{P}$  is a conjunction of predicates that is evaluated for a time series  $\mathcal{V} = \langle v_0, \dots, v_{n-1} \rangle$  of supercap voltages. We have identified three different policies that we introduce and discuss in the following.

1) *Depletion Safety*: Achieving perpetual operation requires depletion-safe operation, hence the operating voltage  $V_c$  of the supercap must not fall below the cut-off voltage  $V_{\text{cut}}$ . Due to uncertain energy intake in the future, we define a critical voltage  $V_{\text{crit}} \geq V_{\text{cut}}$  that must never be undercut. The corresponding policy is

$$\mathcal{P}_{\text{DS}}(\mathcal{V}) := \forall v \in \mathcal{V} : v \geq V_{\text{crit}}. \quad (6)$$

2) *Maximum Power Point*: The harvester introduced in Sect. III operates at a better power-point for higher values of  $V_c$ , since for a given value of  $I_n$ , the regulator input  $I_r$  and  $V_c$  behave reciprocally, cf. (1)—i.e.,  $I_r$  decreases for increasing  $V_c$ , so that a larger portion of  $I_h$  charges the supercap. The power point is thus best for  $V_c = V_{\text{max}}$  and it is desirable to operate the supercap at a high voltage. Requiring a high voltage  $V_c$  at all times yet implies a low load  $I_n$ , so that we only require  $V_c$  to reach  $V_{\text{max}}$  (at least) once a day. The corresponding policy is

$$\mathcal{P}_{\text{MPP}}(\mathcal{V}) := \exists v \in \mathcal{V} : v = V_{\text{max}} \wedge \mathcal{P}_{\text{DS}}(\mathcal{V}). \quad (7)$$

3) *Balanced Consumption*: Based on the principle of energy-neutrality [4], we deduce the balanced-consumption predicate, i.e., a node must not consume more energy than its harvester produces within the prediction horizon.

$$\mathcal{P}_{\text{Bal}}(\mathcal{V}) := v_0 = v_{|\mathcal{V}|-1} \wedge \mathcal{P}_{\text{DS}}(\mathcal{V}). \quad (8)$$

In particular, the final voltage must be the same as the initial voltage; intermediate voltages are not considered.

### B. Finding the Maximum Compliant Load

Given an energy policy  $\mathcal{P}$ , we target at finding the maximum compliant load  $I_n^*$ , so that the resulting time series of voltages  $\mathcal{V}$  satisfies  $\mathcal{P}$ . Here,  $I_n^*$  is constant w.r.t. the prediction horizon to achieve stable, uniform operation and represents an average consumption.

For a given set of parameters, we use the algorithm from Fig. 2 to calculate a series of intermediate voltages  $\mathcal{V}$ ,

```

function SIMVC( $V_c, \Delta t, I_n, I_h$ )
  if  $V_c < V_{\text{cut}}$  then
     $I_n \leftarrow 0$ 
   $y \leftarrow V_c$ 
  if  $I_h = 0$  then
     $a \leftarrow (V_n I_n) / (C \eta)$ 
     $y \leftarrow \sqrt{y^2 - 2a \cdot \Delta t}$ 
  else if  $y \neq a/b$  then
    repeat
       $y_{\text{old}} \leftarrow y$ 
       $y \leftarrow g(y, V_c, \Delta t)$ 
    until  $|y_{\text{old}} - y| < \epsilon$ 
     $y \leftarrow \min(y, V_{\text{max}})$ 
  return  $y$ 

function FINDMAXLOAD( $V_c, \mathcal{S}$ )
   $l \leftarrow 0$ 
   $h \leftarrow I_n^{\text{max}}$ 
  repeat
     $m \leftarrow (h + l) / 2$ 
     $v \leftarrow V_c$ 
     $\mathcal{V} \leftarrow \langle v \rangle$ 
    for all  $(l, \mu) \in \mathcal{S}$  do
       $v \leftarrow \text{SIMVC}(v, l, m, \mu)$ 
       $\mathcal{V} \leftarrow \mathcal{V} \cup \langle v \rangle$ 
    if  $\mathcal{P}(\mathcal{V})$  then
       $l \leftarrow m$ 
    else
       $h \leftarrow m$ 
  until  $h - l < \text{tol}$ 
  return  $l$ 

```

Figure 2. Left: algorithm for supercap voltage simulation; right: algorithm to find the maximum supported load  $I_n^*$  complying with policy  $\mathcal{P}$

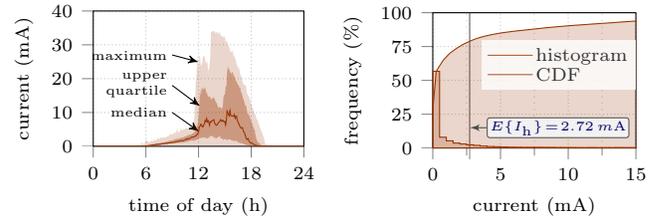


Figure 3. Characteristics of solar current  $I_h$  produced by the harvester

where the individual elements  $v \in \mathcal{V}$  are the voltages at the beginning (or ends, respectively) of time slots. We evaluate  $\mathcal{P}$  w.r.t.  $\mathcal{V}$  to decide whether  $\mathcal{P}$  holds or is violated. Finding the maximum supported load  $I_n^*$  is achieved efficiently by binary search. The maximum number of iterations is limited by an absolute tolerance  $\text{tol}$ , so that the binary search has  $\log_2(I_n^{\text{max}}/\text{tol})$  iteration steps, where  $I_n^{\text{max}}$  is the maximum allowed consumption. The final algorithm is shown in Fig. 2.

## VI. EVALUATION

### A. Data Basis

We analyze the proposed energy policies with a 194-day harvesting trace collected with our harvester, which was placed on a window sill facing westwards. The current  $I_h$  produced by the solar cell was sampled every 30 s, from which we obtained 5 min averages. The daily distribution of  $I_h$  is shown in Fig. 3 (left) and reveals a very distinct course, which is due to shades produced by buildings and a roof overhang. Fig. 3 (right) shows the overall distribution of  $I_h$ . The mean current production is 2.72 mA and almost 80% of all 5 min averages do not exceed the mean. These findings stress the need for online duty-cycle adaptation.

### B. Methodology, Parameters, and Metrics

We simulate the voltage course for an Iris node powered by our harvester with a time resolution of 5 min using our  $I_h$ -trace. The node's consumption  $I_n = I_n^*$  is determined by FINDMAXLOAD for a prediction horizon of 24 h and updated periodically. We use  $I_n^{\text{max}} = 25 \text{ mA}$  and  $\eta = 1$  to

enable a straight-forward comparison of the values  $I_h$  and  $I_n^*$ . This setup complies with the current consumption of an Iris node with enabled radio and the actual average efficiency of our harvester (ca. 20 mA with 80% efficiency).

The supercaps are dimensioned from 25 F to 200 F. We chose  $\epsilon = 0.1$  mV and  $\text{tol} = 10$   $\mu$ A. We test critical voltages  $V_{\text{crit}}$  from 0.5 V to 1.5 V. Forecasts are created with a per-slot EWMA filter using 6 to 96 slots (of equal length). We do not use the WCMA-approach, as it cannot produce long-term predictions, i.e., predictions longer than one slot. To evaluate the policies without bias of the chosen predictor, we generate perfect predictions with and without slots, i.e., we feed the actual  $I_h$ -data of the next 24 h to FINDMAXLOAD.

We assess the performance of the policies in terms of the fractional downtime of the node and the distribution of  $I_n^*$ . The latter is displayed with box plots showing the median along with minimum and maximum values plus upper and lower quartiles. The horizontal gray line in all box plots indicates the average value of  $I_n^*$ .

### C. Results

1) *Perfect Prediction:* The theoretical potential and conceptual weaknesses of the policies are revealed by providing a perfect prediction. Here, the node updates its load  $I_n^*$  every 5 min, i.e., with the granularity of the  $I_h$ -trace and voltage simulation. Hence, no downtimes are experienced. Figure 4 shows the distribution of  $I_n^*$ . The figure reveals that all medians are considerably below this average value. For a 25 F supercap, only one fourth of the harvested current is effectively used by the node. The influence of  $V_{\text{crit}}$  depends on  $C$  (see Fig. 4a): In case of small  $C$ , a larger value of  $V_{\text{crit}}$  leads to a smaller load. Combining  $V_{\text{crit}} = 1.5$  V with a 200 F supercap increases the average load, because the supercap is operated at higher voltages and thus kept in a better (harvesting) power point. This effect is amplified by the fact that a larger  $V_{\text{crit}}$  increases the likelihood of  $V_c$  reaching  $V_{\text{max}}$ . Yet, a too large value of  $V_{\text{crit}}$  will produce the contrary result: Although the power point is good, only a small fraction of energy can be taken from the supercap.

The maximum-power-point policy  $\mathcal{P}_{\text{MPP}}$  brings improvements over plain  $\mathcal{P}_{\text{DS}}$  for  $C \geq 100$  F (see Fig. 4b). This supports the previous findings and shows that small supercaps are fully charged once per day in our setup, even without applying  $\mathcal{P}_{\text{MPP}}$ . Figure 4c indicates that  $\mathcal{P}_{\text{Bal}}$  improves over  $\mathcal{P}_{\text{DS}}$  only for large capacities. Its overall performance does not achieve the one produced by  $\mathcal{P}_{\text{MPP}}$ .

A detailed trace of  $V_c$  and  $I_n$  is shown in Fig. 5. On days with high energy intake,  $\mathcal{P}_{\text{MPP}}$  and  $\mathcal{P}_{\text{DS}}$  show similar voltages traces. Differences are notable, during days with low intake. Particularly in the first days,  $\mathcal{P}_{\text{MPP}}$  achieves higher loads, because  $V_c$  is held at a higher level. The price of this benefit is paid, when a good day is followed by a poor one (e.g., day 13): The load is suddenly decreased, whereas  $\mathcal{P}_{\text{DS}}$  allows for a smoother descent. The main problem of

$\mathcal{P}_{\text{Bal}}$  is evident on days 13 and 15: Due to a drop in intake prediction, the policy is only satisfied for  $I_n = 0$ . Ironically, this situation always arises after particularly good days.

2) *Perfect Prediction with Slots:* We repeated the previous experiments by introducing slots, i.e., FINDMAXLOAD is only executed at the beginning of a slot and the forecast is less fine-grained (there is one average value for each slot). In case of, e.g., 12 slots,  $I_n$  is updated every 2 h and the forecast consists of twelve 2 h averages of  $I_h$ .

The overall performance of the policies does not change considerably compared to the previous results. In general, using less slots slightly decreases  $I_n^*$ . We omit the distribution plot of  $I_n^*$  due to space constraints. The combination of coarse-grained prediction values and less frequent updates results in non-negligible downtimes for  $V_{\text{crit}} = 0.5$  V. Figure 6 shows that less slots lead to longer downtimes. For  $\mathcal{P}_{\text{DS}}$  (see Fig. 6a), downtimes are longer for larger values of  $C$ , because the charging time is longer. In contrast,  $\mathcal{P}_{\text{MPP}}$  and  $\mathcal{P}_{\text{Bal}}$  (Fig. 6b and 6c) enforce larger values of  $V_c$ , so that less downtimes result from larger capacities.

3) *EWMA Prediction with Slots:* For real-world deployments, EWMA predictions with slots are frequently used. We evaluated our policies for a wide range of smoothing factors, but due to page limitation only discuss the results for 0.8, which yields good results in our setup.

The distribution of  $I_n^*$  depicted in Fig. 7 is similar to the one for perfect prediction, but shows a larger variance and less influence of  $V_{\text{crit}}$ . Only for  $C = 200$  F,  $\mathcal{P}_{\text{Bal}}$  and  $\mathcal{P}_{\text{MPP}}$  outperform  $\mathcal{P}_{\text{DS}}$ . The most stable operation, i.e., least variance, is achieved by  $\mathcal{P}_{\text{MPP}}$ . Prediction errors lead to increased downtimes, which is evident from Fig. 8. Particularly for  $V_{\text{crit}} = 0.5$  V,  $\mathcal{P}_{\text{DS}}$  (Fig. 8a) suffers from downtimes of up to 26%. An increase of  $V_{\text{crit}}$  to 1.0 V decreases downtimes, but cannot protect the nodes from failure, when small capacities, offering less energy reserves for the same  $V_{\text{crit}}$ , are used.

The trace in Fig. 9 helps understanding these results. None of the policies can prevent energy shortage on day 17. The good energetic condition of the previous days is interrupted by a day with almost no harvest, so that the EWMA forecast is too high. This results in a low  $V_c$  but no adaptation of  $I_n$ , since the forecast promises large values of  $I_h$ , i.e.,  $I_h \gg I_n^*$ . Due to the poor energetic conditions, it takes 2 days to reach the required start-up voltage of 1.6 V. Fig. 9 also reveals the general reaction of the policies to prediction errors.  $\mathcal{P}_{\text{MPP}}$  produces peaks of  $I_n^*$ , if  $I_h$  exceeds the forecast. In contrast,  $\mathcal{P}_{\text{Bal}}$  decreases  $I_n$  in the same situation: If the harvest exceeds the forecast,  $V_c$  becomes larger than expected. To comply with  $\mathcal{P}_{\text{Bal}}$ , i.e., retain  $V_c$  at the same time on the next day,  $I_n$  must be reduced, since the forecast is only adjusted slightly due to smoothing. The choice of  $V_{\text{crit}}$  only plays a role for  $\mathcal{P}_{\text{MPP}}$  and  $\mathcal{P}_{\text{Bal}}$ , if the forecast promises a good harvest. In subsequently improving conditions (days 9 to 13), a higher but less stable load is achieved in comparison to the

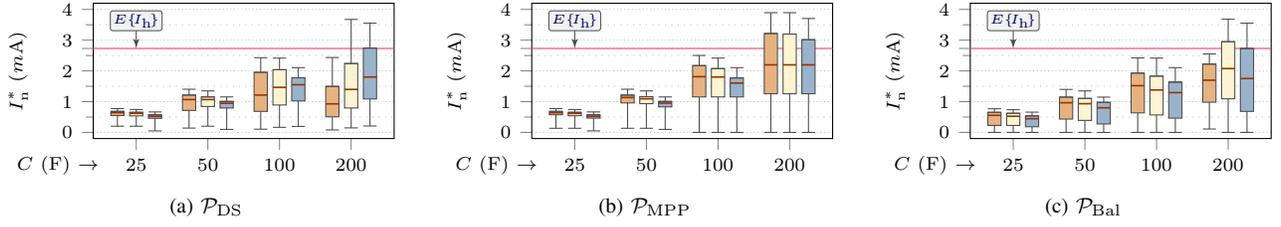


Figure 4. Perfect prediction: Distribution of  $I_n^*$  for  $V_{crit}$ -values of 0.5 V (orange/left), 1.0 V (yellow/center), 1.5 V (blue/right)

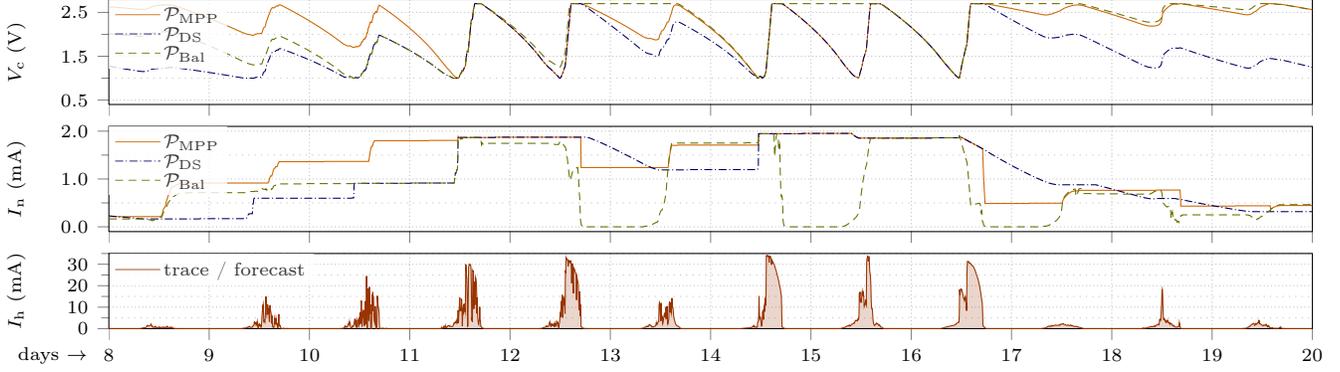


Figure 5. Perfect prediction: Course of  $V_c$  and  $I_n$  for  $V_{crit} = 1.0$  V and  $C = 100$  F. Ticks on the x-axis indicate midnight

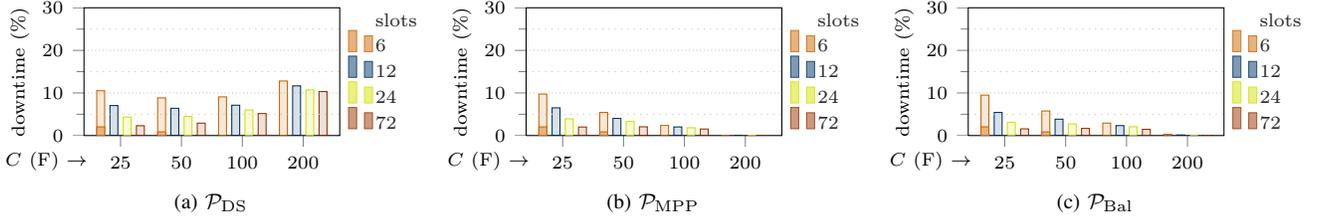


Figure 6. Perfect prediction with 6 to 72 slots: Influence of capacity  $C$ , policy  $\mathcal{P}$  on node downtime for  $V_{crit} = 0.5$  V (light) and  $V_{crit} = 1.0$  V (dark)

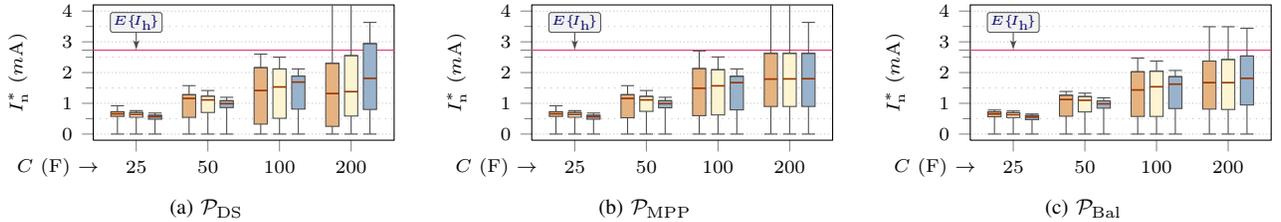


Figure 7. EWMA prediction: Distribution of  $I_n^*$  using 12 slots and an EWMA smoothing factor of 0.8. The critical voltage is  $V_{crit} = 1.0$  V

other policies. The downside of this behavior is a lower  $I_n$  on single poor days (e.g., day 13) and subsequently decreasing harvesting conditions (not shown in the figure).

#### D. Summary and Discussion

The results show that  $\mathcal{P}_{MPP}$  generally yields the highest mean and lowest deviation of  $I_n^*$ . For our prototype, the gap vanishes for small values of  $C$ , as the solar cell charges the supercaps even on days with low harvest. This also causes a relatively low average  $I_n^*$  compared to the mean  $I_h$ , since a surplus  $I_h > I_r$  cannot be buffered for later use.

Step length  $\Delta t$  influences the accuracy of SIMVC notably: The more coarse-grained  $I_h$  and  $I_n$  are, the less accurate  $\mathcal{V}$  becomes. While the impact on the distribution of  $I_n^*$  is rather low, huge downtimes are experienced, if a perfect forecast with only 6 slots meets a value of  $V_{crit}$  close to  $V_{cut}$ .

The results of the EWMA forecasts produce several insights: Forecast errors reduce the improvements of  $\mathcal{P}_{MPP}$ , and ignoring a day's trend provokes depletion. Choosing  $V_{crit} \gg V_{cut}$  guards against depletion while having a small impact on  $I_n^*$  in most cases.  $\mathcal{P}_{Bal}$  does not work in praxis:  $I_n^*$  becomes smaller, when the forecast is exceeded.

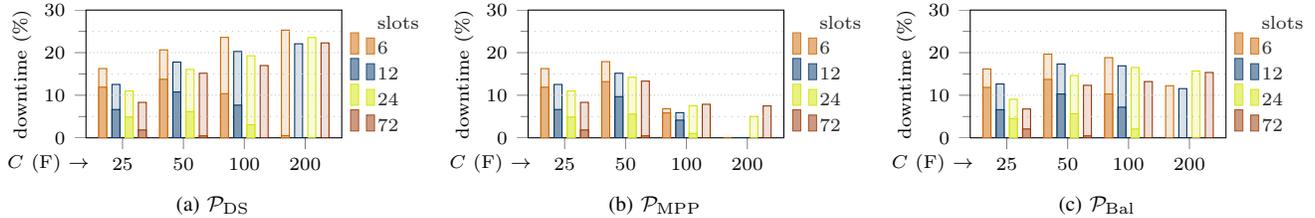


Figure 8. EWMA prediction: Influence of capacity  $C$  and policy  $\mathcal{P}$  on node downtime for  $V_{\text{crit}} = 0.5 \text{ V}$  (light) and  $V_{\text{crit}} = 1.0 \text{ V}$  (dark)

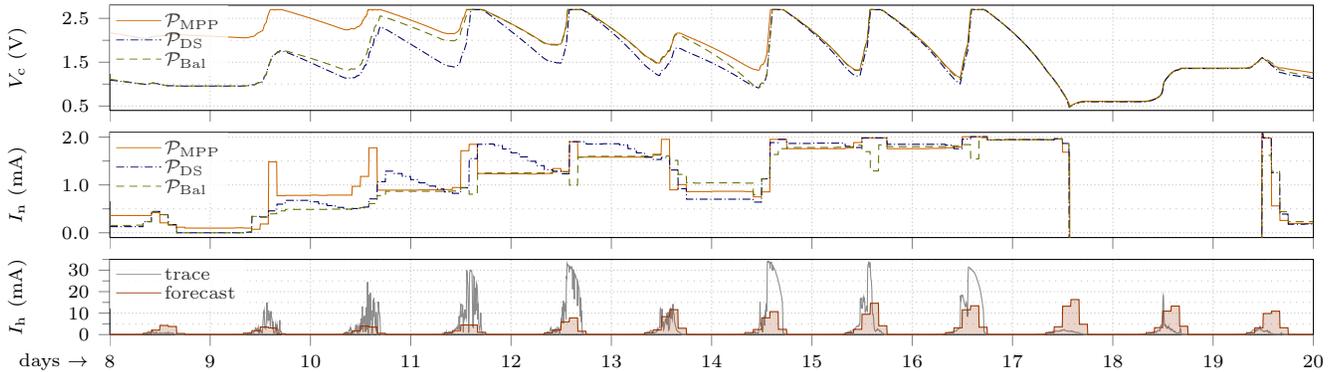


Figure 9. EWMA forecast: Course of  $V_c$  and  $I_n$  for  $V_{\text{crit}} = 1.0 \text{ V}$  and  $C = 100 \text{ F}$  using 12 slots and an EWMA filter coefficient of 0.8

## VII. CONCLUSION

We presented a policy-based concept for online adaptation of an energy-harvesting sensor node’s consumption. Our approach is tailored to (non-linear) systems solely relying on supercaps as energy buffer, for which existing solutions are inappropriate. We devised three policies to achieve depletion-safe and uniform, utility-maximal node operation, which we evaluated using a real solar-harvest trace.

The evaluation shows the applicability of the policies and confirms that operating a supercap close to its maximum voltage yields improved mean utility with low variation. The results stress that choosing a duty cycle or task schedule a priori is likely to either waste energetic resources or lead to unacceptable node downtimes. Our policy concept achieves rapid adaptation to the expected energy harvest. Enforcing a minimum voltage slightly above the harvester’s cut-off voltage decreases the risk of depletion but does not affect utility. This is particularly important for forecasts using time slots: Prediction errors and model inaccuracy caused by the linearization due to long time slots must be compensated.

We also identified shortcomings. In a real deployment, depletion cannot be always prevented. More precise energy-harvest forecasts are needed to increase the performance of the policy concept. An approach similar to WCMA is an option, but needs detailed analysis. A second one is pessimistic forecasting, e.g., predicting a lower bound of energy intake per slot, but it requires a strategy for coping with excess energy. Our current approach lacks an explicit feedback loop, such as exploiting the error of predicted versus real voltage to detect changing harvest conditions.

## REFERENCES

- [1] X. Jiang, J. Polastre, and D. Culler, “Perpetual Environmentally Powered Sensor Networks,” in *IPSN*, 2005.
- [2] F. Simjee and H. C. P. “Everlast: Long-Life, Supercapacitor-Operated Wireless Sensor Node,” in *ISLPED*, 2006.
- [3] C. Moser, L. Thiele, D. Brunelli, and L. Benini, “Robust and Low Complexity Rate Control for Solar Powered Sensors,” in *DATE*, 2008.
- [4] A. Kansal, J. Hsu, S. Zahedi, and B. S. M. “Powermanagement in Energy Harvesting Sensor Networks,” *Trans. on Embedded Computing Sys.*, 2007.
- [5] M. Vigorito, C. D. Ganesan, and G. Barto, A. “Adaptive Control of Duty Cycling in Energy-Harvesting Wireless Sensor Networks,” in *SECON*, 2007.
- [6] C. Renner, J. Jessen, and V. Turau, “Lifetime Prediction for Supercapacitor-powered Wireless Sensor Nodes,” in *FGSN*, 2009.
- [7] A. Weddell, G. Merrett, T. Kazmierski, and B. Al-Hashimi, “Accurate Supercapacitor Modeling for Energy-Harvesting Wireless Sensor Nodes,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, Sep. 2011.
- [8] C. Renner, S. Unterschtz, and V. Turau, “Power Management for Wireless Sensor Networks Based on Energy Budgets,” Hamburg University of Technology, Hamburg, Germany, Tech. Rep., Jul. 2011.
- [9] P. Dutta, J. Hui, J. Jeong, S. Kim, C. Sharp, J. Taneja, G. Tolle, K. Whitehouse, and D. Culler, “Trio: Enabling Sustainable and Scalable Outdoor Wireless Sensor Network Deployments,” in *IPSN*, 2006.
- [10] J. Recas Piorno, C. Bergonzini, D. Atienza, and T. Simunic Rosing, “Prediction and Management in Energy Harvested Wireless Sensor Nodes,” in *VITAE*, 2009.
- [11] M. Ali, B. Al-Hashimi, J. Recas, and D. Atienza, “Evaluation and Design Exploration of Solar Harvested-Energy Prediction Algorithm,” in *DATE*, 2010.