Poster Abstract: Long-term Energy-neutral Operation of Solar Energy-harvesting Sensor Nodes under Time-varying Utility

Kai Geissdoerfer
CSIRO, TU Berlin, TU Dresden
kai.geissdoerfer@tu-dresden.de

Raja Jurdak
CSIRO
raja.jurdak@csiro.au

Brano Kusy
CSIRO
brano.kusy@csiro.au

ABSTRACT

Sensor networks increasingly rely on harvesting energy from the environment to sense, process, and transmit data. Online energy availability forecasting and energy management are critical to ensure long-term energy-neutral operation of battery-powered energy-harvesting sensor nodes. Existing methods focus on applications with time-invariant utility and custom-tailored hardware platforms, which limits their effectiveness across diverse application domains, different platforms, and in the face of aging hardware components. To address these limitations, we formulate an optimisation problem with respect to time-varying utility under the given hardware constraints. We also present *PREACT*, an online energy-management algorithm that approximates the optimal solution to the optimisation problem by incorporating long-term energy forecasting.

1 INTRODUCTION

As wireless sensing systems enter commercial contexts, long-term unattended operation becomes paramount to their success. Given the limited capacity of batteries, this goal can be achieved if the mean energy consumed equals the mean energy harvested (energy neutrality) from sources such as solar radiation or vibrations.

To this end, sensor nodes limit their energy consumption to a given budget by scheduling sampling, processing, or communication accordingly. An energy-management algorithm computes this energy budget based on estimates of current and future energy availability. Previous work has focused on minimising temporal variations of the assigned budget [2, 7]. Many practical applications, however, have a time-varying utility. For example, in sensing applications utility may be the "value" of information extracted from collected data. The temporal trend of utility is often known a priori due to, for example, seasonal or periodic trends in the sensed phenomena. Energy management should take such trends into account by assigning the budget proportionally to the time-varying utility. In this way, the overall benefit of a sensing system can be improved by conserving energy when the information content of the data is low, so that more energy is available during the high-utility periods.

Existing methods for energy management can be classified into predictive and reactive. Predictive methods forecast energy availability to strategically assign a budget according to given performance criteria. Most methods for solar energy-harvesting nodes focus on managing energy over short time intervals (i.e., days or weeks). Solar radiation, however, has strong seasonal variations, which should be considered to ensure long-term operation of a sensing system. The only long-term methods [1, 2] rely on a well-dimensioned system according to a capacity-planning method based on knowledge of application power requirements and deployment location [3]. Achieving long-term energy-neutral operation is then a relatively simple task. However, solar panel size and battery weight can be constrained in practice, such as when tracking small assets or animal species. Perhaps more importantly, batteries and solar panels are subject to aging.

Reactive energy-management methods adapt to deviations of the battery's actual state of charge (SOC) from a target SOC by changing the system's energy spending [6, 7]. This approach, however, cannot prepare for future energy availability, for example, by hoarding energy during high-availability periods to be able to sustainably deliver performance during times of low energy availability.

To tackle the limitations of prior work, we present *PREACT*, an online algorithm for energy management of sensor nodes with solar energy harvesting under time-varying utility. *PREACT* is a hybrid solution, combining predictive and reactive elements. It computes a target SOC based on predicted long-term solar energy availability and reactively tracks this target SOC using a control loop.

2 OPTIMISATION PROBLEM

The user formally specifies application utility as a function of time with values ranging between 0 and 1, where 1 indicates maximum utility. The ideal energy budget is a linearly scaled version of the utility function. This linear scaling translates the utility into an actual energy budget that fulfils the energy-neutrality condition.

The amount of energy that must be temporarily stored in the battery depends crucially on the coherence of harvested energy and utility function. The ideal energy budget is not sustainable if the battery capacity is insufficient. Furthermore, an online energy-management algorithm will likely not be able to perfectly follow the ideal budget. The deviation of the assigned energy budget from the ideal energy budget is a metric for the underperformance of an energy management algorithm relative to the theoretical optimum.

We informally state the energy-management problem as a constrained optimisation problem as follows: *Find the energy budget that minimises underperformance, such that all incoming energy is utilised and the SOC never exceeds the capacity of the battery.*

3 PREACT

PREACT is an online algorithm that approximates the optimal solution to the above optimisation problem. It targets energy management on a single sensor node with solar harvesting capability.

PREACT executes with a given period (e.g., every day). Based on the harvested energy during the preceding period and the current SOC, it computes the optimised budget and resulting SOC for one year ahead (to account for the variation in solar radiation throughout the year). The computed budget for the following period is then

0-7695-6377-5/18/$31.00 ©2018 IEEE
DOI 10.1109/IPSN.2018.00041

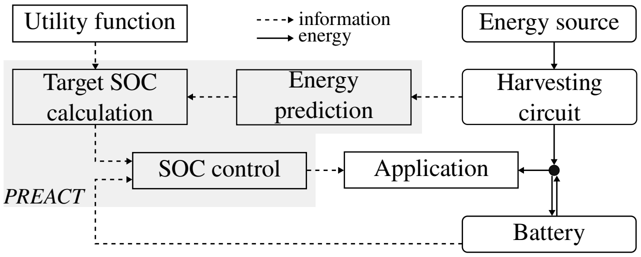

Figure 1: Using predictions of energy availability and a user-defined utility function, *PREACT* periodically determines a target SOC. It uses a controller to follow this target SOC, periodically assigning an optimised budget to the application.

assigned to the system. Figure 1 gives an overview of *PREACT*. It consists of three main components that work in a cascade.

Energy prediction. This component provides one-year-ahead predictions of harvested energy, covering relevant seasonal variations of solar energy. We extend a popular astronomical model [3] to be hardware and location agnostic by learning the model parameters online. With each measurement of harvested energy over the preceding period, *PREACT* updates the model parameters towards the negative gradient of the mean squared error (MSE) between the predicted energy input and the actual harvested energy.

Target SOC calculation. The time-dependent target SOC is calculated based on the utility function and the energy prediction. The trend of SOC that would result from spending the ideal budget is calculated by considering the predicted energy inputs and planned outputs. If the calculated SOC exceeds the capacity of the battery, the planned budget has to be adjusted to spend excess energy that cannot be stored in the battery. This adjustment is implemented by linear scaling of the target SOC, trading off precise adherence to the ideal energy budget against feasibility to realise it.

SOC control. *PREACT* uses a proportional-integral-derivative (PID) controller to follow the target SOC by assigning an energy budget for the following period. The PID controller allows for a lightweight implementation and proved useful for SOC control [6].

4 PRELIMINARY EVALUATION

We use trace-based simulations to compare *PREACT* against two state-of-the-art methods, Short-term Exponentially Weighted Moving Average (*ST-EWMA*) [5] and Long-term Energy Neutral Operation (*LT-ENO*) [2], in terms of underperformance.

Method. We consider an energy-harvesting node with a 5000 mAh battery and a 3500 mm^2 solar panel. The simulation model covers real-world effects, including charging efficiencies, capacity degradation, self-discharge, and SOC estimation error. We feed our simulations with satellite-derived solar energy traces for years 1989 to 1994 from five locations in all major climate zones. These data were obtained from the NASA Langley Research Center Atmospheric Science Data Center Surface meteorological and Solar Energy (SSE) web portal supported by the NASA LaRC POWER Project [4]. All algorithms execute with a period of one day in our simulations.

Results. Figure 2 plots underperformance for three different utility functions. We see that *PREACT* improves underperformance by 40 % and 50 % relative to *LT-ENO* and *ST-EWMA*, respectively, across all utility functions. *LT-ENO* performs poorly for time-varying utility, and falls behind *PREACT* even for uniform utility as it assumes a specifically dimensioned system without aging effects. *ST-EWMA* closely couples utility with energy availability, leading to large underperformance during times of high demand and low supply.

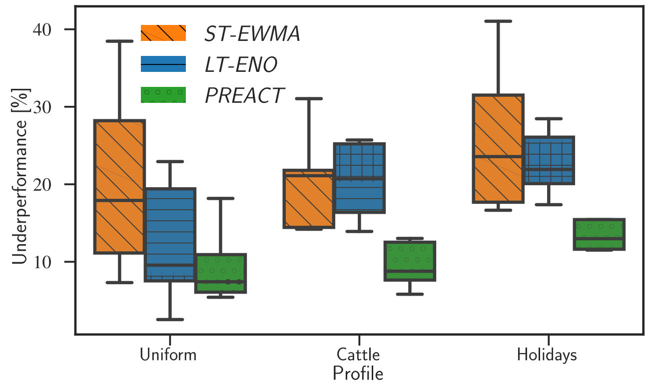

Figure 2: Underperformance for three utility functions. The *Cattle* profile exhibits a seasonal component, where utility of GPS tracking is lower, while cattle are locked in small paddocks during the rain season. The high frequency *Holidays* profile assumes higher utility for beach water monitoring on weekends and holidays, when beaches are busy.

5 CONCLUSIONS

We have explored how incorporating prior knowledge about the utility of an application into energy management can help optimise performance. *PREACT* is the first energy-management algorithm that works on arbitrarily dimensioned systems independent of the deployment location. It uses a simple optimisation algorithm to trade energy utilisation for adherence to the utility function. Our results show that *PREACT* outperforms the state of the art.

Acknowledgements. This work was partially supported by DFG within the Cluster of Excellence CFAED and by the DAAD scholarship programme PROMOS. We thank Marco Zimmerling for his valuable feedback received on an early version of this paper.

REFERENCES

[1] B. Buchli, P. Kumar, and L. Thiele. 2015. Optimal power management with guaranteed minimum energy utilization for solar energy harvesting systems. In *IEEE DCOSS*.
[2] B. Buchli, F. Sutton, J. Beutel, and L. Thiele. 2014. Dynamic Power Management for Long-term Energy Neutral Operation of Solar Energy Harvesting Systems. In *ACM SenSys*.
[3] B. Buchli, F. Sutton, J. Beutel, and L. Thiele. 2014. Towards Enabling Uninterrupted Long-Term Operation of Solar Energy Harvesting Embedded Systems. In *EWSN*.
[4] NASA Atmospheric Science Data Center. 2016. Surface meteorology and Solar Energy. https://eosweb.larc.nasa.gov/sse/. (2016).
[5] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava. 2007. Power Management in Energy Harvesting Sensor Networks. *ACM Transactions on Embedded Computing Systems* 6, 4 (2007).
[6] T. N. Le, O. Sentieys, O. Berder, A. Pegatoquet, and C. Belleudy. 2012. Power manager with PID controller in energy harvesting wireless sensor networks. In *IEEE GreenCom*.
[7] C. M. Vigorito, D. Ganesan, and A. G. Barto. 2007. Adaptive Control of Duty Cycling in Energy-Harvesting Wireless Sensor Nodes. In *IEEE SECON*.