

Diss. ETH No. 22717

Power Subsystem Design and Management for Solar Energy Harvesting Embedded Systems

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

Bernhard Raget Buchli

M.Sc. EE, University of Wisconsin-Madison, USA

born on June 13, 1976

citizen of Safien, GR

accepted on the recommendation of

Prof. Dr. Lothar Thiele, examiner

Prof. Dr. Luca Benini, co-examiner

Dr. Jan Beutel, co-examiner

2015



Institut für Technische Informatik und Kommunikationsnetze
Computer Engineering and Networks Laboratory

TIK-SCHRIFTENREIHE NR. 153

Bernhard Raget Buchli

Power Subsystem Design and Management for Solar Energy Harvesting Embedded Systems



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

A dissertation submitted to
ETH Zurich
for the degree of Doctor of Sciences

Diss. ETH No. 22717

Prof. Dr. Lothar Thiele, examiner
Prof. Dr. Luca Benini, co-examiner
Dr. Jan Beutel, co-examiner

Examination date: May 27, 2015

For D. L. A. S.,
who makes my life complete!

Abstract

Wireless Sensor Networks (WSN) have reached a level of maturity, at which they have become a feasible option for monitoring processes of interest, when wired infrastructure is not possible. Due to remote and inaccessible deployment sites of such networks, they can generally not rely on reliable power sources, but require batteries to supply the energy for the system to perform its intended task. However, the finite energy store imposed by batteries directly limits the system's performance and lifetime. Ambient energy harvesting has been shown to be a promising way to boost the performance and lifetime of WSNs. Unfortunately, enhancing a battery powered device with energy harvesting capabilities will by itself neither provide a lower bound on the expected sustainable performance level, nor guarantee uninterrupted long-term operation.

This thesis addresses the design and runtime management of the power subsystem for solar energy harvesting embedded systems. We demonstrate that for enabling long-term operation of such systems, a paradigm shift both in the design approach, and the runtime management of the energy is necessary. We provide an end-to-end power management solution, which consists of (i) a power subsystem capacity planning approach, and (ii) two novel dynamic power management schemes that maximize the minimum achievable performance level, while ensuring that long-term, i.e., multi-year operation can be sustained. Compared to three State-of-the-Art approaches, our solution maximizes the long-term sustainable minimum system performance or perform equivalently, but require a smaller solar panel and/or smaller battery. Our theoretical results are supported by simulations using 10 years of solar energy measurements from various geographical locations. To demonstrate the improvements of proper power subsystem design and management, we further present a case study with a real-world WSN deployment for geoscientific research in a high-alpine environment.

Specifically, this thesis presents the following contributions to the State-of-the-Art:

- We present a systematic method for power subsystem capacity planning, i.e., proper sizing of the solar panel and battery, for solar energy harvesting embedded systems.

- We present two novel dynamic power management schemes that enable the system to maximize the minimum performance level at runtime, while ensuring that uninterrupted operation over time periods on the order of years may be sustained.
- We present a light-weight battery State-of-Charge approximation algorithm that can provide runtime information about the battery fill level without requiring special purpose hardware. Moreover, we show that this approach can be used to infer the harvested energy.
- Through extensive simulation we show that our end-to-end solution achieves significantly better results in terms of minimum long-term sustainable performance level, duty-cycle stability, and overall energy efficiency when compared to 3 State-of-the-Art approaches.
- Finally, using a real-world scientific project, we demonstrate the improvements in system utility for the end-user application that can be gained with our solution.

Zusammenfassung

Drahtlose Sensornetzwerke (WSN) haben einen hohen Reifegrad erreicht, so dass sie eine praktikable Lösung für die Überwachung von verschiedenen Prozessen sind, wenn verdrahtete Infrastruktur nicht möglich ist. Aufgrund der oft schwer zugänglichen Einsatzorte solcher Netzwerke können die einzelnen Geräte in der Regel nicht zuverlässig mit Strom versorgt werden, sondern es sind Batterien erforderlich, um die Energie für das System zu liefern. Allerdings begrenzt der endliche Energiespeicher von Batterien die Leistung und Lebensdauer des Systems. Energieentnahme aus der Umgebung (ambient energy harvesting) ist ein vielversprechender Weg, um die Leistung und Lebensdauer des WSNs zu steigern. Leider verbessert die Bestückung eines batteriebetriebenen Gerätes mit Energy-Harvesting-Fähigkeiten weder automatisch das erwartete Leistungsniveau, noch garantiert es ununterbrochenen Langzeitbetrieb.

Diese Dissertation befasst sich mit dem Design und der Laufzeitverwaltung des Stromversorgungssystems für eingebettete Systeme mit Sonnenenergiegewinnungsfähigkeiten. Es wird gezeigt, dass zur Ermöglichung von langfristigem Betrieb dieser Systeme ein Paradigmenwechsel sowohl für den Designansatz als auch für die Laufzeitverwaltung der verfügbaren Energie notwendig ist. Es wird eine komplette Power-Management-Lösung vorgestellt, welche aus einem Kapazitätsplanungsansatz des Energieversorgungssystems und aus zwei neuen dynamischen Power-Management-Algorithmen besteht, welche dafür sorgen, das minimale Leistungsniveau zu maximieren. Sie stellen gleichzeitig sicher, dass langfristiger, das heißt, mehrjähriger Betrieb aufrechterhalten werden kann.

Im Vergleich zum letzten Stand der Technik, maximiert unsere Lösung die langfristige nachhaltige Mindestleistung des Systems oder erreicht mindestens äquivalente Performanz, aber mit einer kleineren Solarzelle und/oder kleineren Batterie. Unsere theoretischen Ergebnisse werden durch Simulationen mit Sonnenenergiemessungen während 10 Jahren an verschiedenen geografischen Standorten unterstützt. Um die Nutzungssteigerung unseres Ansatzes zu demonstrieren, evaluierten wir unsere Lösung anhand einer Fallstudie mit einer realen WSN-Installation für die geowissenschaftliche Forschung in einer hochalpinen Umgebung.

Diese Dissertation macht die folgenden Beiträge zum letzten Stand der Technik:

- Wir präsentieren eine systematische Methode für die Kapazitätsplanung des Stromversorgungs-System, das heisst, die richtige Dimensionierung der Solarzelle und Batterie für eingebettete System mit Sonnenenergiegewinnung.
- Wir präsentieren zwei neue dynamische Power-Management-Algorithmen, die dem System ermöglichen das minimale Leistungsniveau während der Laufzeit zu maximieren und gleichzeitig sicherzustellen, dass ein ununterbrochener Betrieb über Zeiträume in der Größenordnung von mehreren Jahren aufrechterhalten werden kann.
- Wir präsentieren einen Batteriefüllstand Approximations-Algorithmus, welcher während der Laufzeit, ohne auf spezielle Hardware angewiesen zu sein, Informationen über den Batteriefüllstand zur Verfügung stellt. Darüber hinaus zeigen wir, dass mit diesem Ansatz auch die gewonnene Energie approximiert werden kann.
- Durch umfangreiche Simulationen zeigen wir, dass unsere Lösung deutlich bessere Ergebnisse in Bezug auf das minimale, und langfristig aufrechterhaltene Leistungsniveau, die Stabilität des Auslastungsgrades und der gesamten Energieeffizienz erzielt als 3 Ansätze aus der Literatur.
- Mit einem realen, wissenschaftlichen Projekt zeigen wir die Verbesserungen in der Nützlichkeit des Systems für die Endnutzeranwendung, welche mit unserer Lösung gewonnen werden können.

Acknowledgement

First of all, I would like to express my gratitude to Prof. Dr. Lothar Thiele for giving me the opportunity to write this thesis in his research group. Thank you very much for your guidance, support, and inspiring advice throughout the years.

I would further like to thank my co-examiners Prof. Dr. Luca Benini and Dr. Jan Beutel for their willingness to review this thesis and to serve on the committee board.

This work would have not been possible without the many fruitful discussions with my office mate Felix Sutton, and his significant contributions to the work presented in this thesis. I would also like to extend my gratitude to all current and former colleagues at the Computer Engineering and Networks group at ETH Zurich.

Last but not least, I would like to express my gratitude to my family for supporting me during the last years. I am indebted and grateful to my parents for their constant support and encouragement that I received throughout my studies. Most importantly, the biggest thank you goes to my wife for her love, immense patience, and tireless support!

The work presented in this thesis was scientifically evaluated by the SNSF and financed by the Swiss Confederation and by Nano-Tera.ch. This support is gratefully acknowledged.

Contents

Abstract	iii
Zusammenfassung	v
Acknowledgement	vii
List of Figures	xiii
List of Tables	xv
1 Introduction	1
1.1 Wireless Sensor Networks	1
1.2 Energy Harvesting in WSNs	3
1.2.1 Ambient Energy Harvesting	4
1.2.2 Benefits of Solar Energy Harvesting	5
1.2.3 Challenges in Design and Operation of Solar Energy Harvesting Systems	6
1.3 Review of Power Management Strategies	10
1.3.1 Capacity Planning	10
1.3.2 Dynamic Power Management	12
1.4 Thesis Outline and Contributions	14

2	Power Subsystem Capacity Planning	21
2.1	Introduction	21
2.2	Concept	24
2.3	System Model	25
2.3.1	Energy Availability Model	26
2.3.2	Harvesting Conditioned Energy Model	27
2.4	Capacity Planning Algorithm	31
2.5	Evaluation	32
2.5.1	Evaluation Methodology, Validation Data, and Performance Metrics	33
2.5.2	Energy Availability Model Validation	34
2.5.3	Capacity Planning Performance Evaluation	35
2.6	Sensitivity Analysis	40
2.6.1	Environmental Parameter	42
2.6.2	Setup Parameters	42
2.7	Chapter Summary	44
3	Dynamic Power Management for Long-Term Operation	45
3.1	Introduction	45
3.2	Capacity Planning for Long-Term Energy Neutral Operation	48
3.2.1	System Architecture, Load Model, and System Utility	50
3.2.2	Harvesting Conditioned Energy Availability Model	50
3.2.3	Power Subsystem Dimensioning	52
3.3	Dynamic Power Management	53
3.3.1	Dynamic Performance Scaling	53
3.3.2	Worst Case Energy Conditions	57
3.4	Evaluation	58
3.4.1	Experimental Setup	58
3.4.2	Experimental Results	61
3.5	Benefits of Dynamic Power Management	65
3.6	Sensitivity Analysis	67
3.6.1	Energy Profile Periodicity	67
3.6.2	Panel Inclination and Orientation	68
3.6.3	History Window Size	68
3.7	Practical Considerations	69
3.7.1	Measurement Support	69
3.7.2	Global Time Knowledge	69
3.7.3	Battery Inefficiencies	70
3.7.4	Algorithm Considerations	70
3.8	Chapter Summary	71

4	Optimal Power Management	73
4.1	Introduction	73
4.2	Problem Definition	76
4.2.1	System Model	77
4.2.2	Optimality Criteria of the Controller	78
4.3	Optimal Control	79
4.3.1	Necessary Conditions for Optimality	79
4.3.2	Algorithm	83
4.4	Finite Horizon Control	85
4.4.1	Optimal Periodic Control	85
4.4.2	Finite Horizon Scheme	86
4.4.3	Guarantees	87
4.5	Implementation	89
4.6	Algorithm Parameterization	91
4.6.1	Energy Estimation Model	91
4.6.2	Horizon of Control	91
4.6.3	Frequency of Control	92
4.7	Evaluation	93
4.7.1	Experimental Setup	93
4.7.2	Experimental Results	96
4.7.3	Practical Considerations: Energy Estimation	102
4.8	Chapter Summary	104
5	State-of-Charge Approximation	105
5.1	Introduction	105
5.2	State-of-Charge Approximation Methods	107
5.3	Concept	110
5.4	System Architecture and Assumptions	111
5.5	Battery Model	112
5.5.1	Capacity and Cut-off Voltage	113
5.5.2	Charge and Discharge Profiles	114
5.5.3	Battery Discharging Model	115
5.5.4	Battery Charging Model	118
5.5.5	Model Parameterization	121
5.6	Online State-of-Charge Approximation	124
5.6.1	Battery State Tracking	124
5.6.2	Voltage Measurements	127
5.7	Evaluation	128
5.7.1	Experimental Set-up	128
5.7.2	Experimental Results	129
5.8	Measurement Considerations	138
5.9	Chapter Summary	139

6	Case Study: X-Sense GPS System	141
6.1	Introduction	141
6.2	System Model	143
6.2.1	System Architecture	143
6.2.2	System Utility	146
6.3	Implementation Details	147
6.3.1	Circumventing Limited Measurement Support . . .	147
6.3.2	LT-ENO Algorithm Modification	148
6.4	Performance Evaluation	149
6.4.1	Evaluation Results	149
6.5	Chapter Summary	158
7	Conclusions	161
7.1	Main Results	161
7.2	Possible Future Perspectives	163
	List of Acronyms	167
	Bibliography	167
	List of Publications	181
	Curriculum Vitae	185

List of Figures

1.1	Battery powered Wireless Sensor Network (WSN) motes.	4
1.2	A solar energy harvesting WSN mote.	5
1.3	Power profile for 4 days in January and June.	7
1.4	Short- and long-term evolution of harvested energy.	8
1.5	Chapter structure of this thesis.	14
1.6	System architecture assumed throughout this thesis.	15
2.1	Design flow for power subsystem capacity planning.	24
2.2	Energy model for capacity planning	31
2.3	Ratio of predicted energy versus true energy.	35
2.4	Ratio of total energy approximated, and effectively harvested.	41
2.5	Performance level achieved with 3 different datasets.	41
2.6	Illustration of incident energy with varying inclination and azimuth angles.	43
3.1	Process flow for long-term solar energy harvesting capacity planning and dynamic power management.	49
3.2	Exemplary solar energy profile.	51
3.3	Dynamic power management algorithm example.	55
3.4	Summary of dynamic power management evaluation results.	63
3.5	Benefits of dynamic power management.	66
3.6	Maximum, mean, and minimum achievable duty-cycle.	67
3.7	Sensitivity analysis of history window size W	69
4.1	Illustration of optimal use function.	82
4.2	Minimum utilization and utility.	92
4.3	Optimal use function.	97
4.4	Piecewise linear correction function.	97
4.5	Use function with optimal LUT implementation.	98
4.6	Use function computed with LT-ENO.	99
5.1	System architecture assumed in this chapter.	111
5.2	Manufacturer provided Depth-of-Discharge versus terminal voltage profile.	112

5.3	Relation between terminal voltage and discharge current. .	114
5.4	Qualitative illustration of discharge traces.	116
5.5	Interpolated Depth-of-Discharge approximation.	118
5.6	Qualitative illustration of bulk charging approximation. . .	120
5.7	Discharge traces and resulting approximation errors. . . .	122
5.8	Linear approximation of bulk charging.	124
5.9	State Machine for charge/discharge tracking.	126
5.10	Simplified model of the voltages and currents of the system platform.	128
5.11	Discharge test at differing temperatures.	131
5.12	Discharge test with aged and new batteries.	133
5.13	Discharge test with bimodal load.	135
5.14	Runtime prediction.	136
6.1	System architecture assumed in this chapter.	142
6.2	Error performance and system utility for varying duty-cycles.	147
6.3	LT-ENO performance results for the DH location.	151
6.4	LT-ENO performance results for the GG location.	153
6.5	Comparison of static vs. dynamic approaches.	154
6.6	OPT performance results for DH location.	155
6.7	OPT performance results for GG location.	157

List of Tables

2.1	Datasets used for evaluation of the capacity planning approach.	33
2.2	Energy availability model validation statistics.	36
2.3	Capacity planning evaluation results.	39
3.1	Datasets used for evaluation of the LT-ENO power management algorithm.	59
3.2	Capacity planning results.	60
3.3	Capacity planning evaluation results.	60
3.4	Performance results averaged over all simulation runs. . .	64
4.1	Non-volatile memory requirements for the LUT.	90
4.2	Datasets used for evaluation of the OPT power management algorithm.	95
4.3	Summary of performance results.	103
5.1	Coefficients for State-of-Charge model.	125
5.2	State transition matrix.	126
5.3	Batteries used for experimental validation.	129
5.4	Physical measurement details.	130
5.5	Performance results for the State-of-Charge approximation.	138
6.1	Deployment details for the case study.	146
6.2	Summary of performance results for the case study with LT-ENO.	154
6.3	Summary of performance results for the case study with OPT.	158

1

Introduction

This thesis presents novel techniques for the design and runtime management of the power subsystem for energy harvesting networked embedded systems. While our methods are generally applicable to a wide range of systems that leverage regenerative energy sources, throughout this thesis, we focus on solar powered Wireless Sensor Networks as the primary application scenario, and discuss how the proposed end-to-end power management solution enables dependable, uninterrupted system operation over multiple years despite relying on a highly variable energy source.

In *Section 1.1* we briefly discuss the relevant characteristics of Wireless Sensor Networks, and review common application scenarios. In *Section 1.2* we detail the benefits of energy harvesting in Wireless Sensor Networks scenarios, and elaborate on the new challenges presented by operating such systems over long time periods. Then, in *Section 1.3*, we survey State-of-the-Art energy management strategies. Finally, in *Section 1.4* we present an outline of this thesis, and summarize the key contributions.

1.1 Wireless Sensor Networks

While there exist many different types of applications scenarios, the class of Wireless Sensor Networks (WSNs) that we consider in this thesis consist of a network of autonomously operating, self-powered sensor systems that monitor various processes or events of interest and coopera-

tively communicate the measurements to a central basestation. WSNs are a valid option whenever wired infrastructure is impractical or infeasible. The sensing systems, or so-called mote-class computing devices, are comprised of a low-power micro-controller, internal and/or external, non-volatile memory, a radio chip, and any number of sensors and/or actuators. These sensors may include simple humidity and temperature sensors [HTB⁺08], or may consist of more complex sensing systems like audio/video equipment [GSGSGH11], tilt- and seismometers [WALR⁺06], or GPS receivers [BBF⁺11], among others.

Advances in miniaturization and low-power design of electronic devices have allowed WSNs to reach a state at which they represent a feasible option for long-term continuous observation of various processes at high temporal and spatial resolution. Exemplary WSN deployments include a variety of application scenarios: they range from patient monitoring in health care settings, e.g., [KLS⁺10], over wild-life surveillance, e.g., [ZSLM04, SKMJ14, DEM⁺10], to biological and geological processes monitoring, e.g., [WALR⁺06, HTB⁺08, WBB⁺13, CB10, BBF⁺11], and infrastructure monitoring and failure warning, e.g., [BRVK11, FGJK13, AAA⁺07], to name a few.

While the application scenarios mentioned above target different objectives, the systems employed for their respective purpose exhibit several common requirements. First, the potentially large number of deployed motes and their possibly sacrificial nature, i.e., the motes may not be recoverable once deployed, dictate that the motes are manufactured cheaply. Second, to reduce the intrusiveness in the environment to be monitored, and lessen the deployment cost, the individual mote should be small and lightweight. Finally, and most importantly, the motes are expected to operate reliably, and cooperatively over time periods on the order of months to years under possibly very challenging conditions.

Due to the prevalence of remote or inaccessible deployment sites dictated by many application scenarios, and the resulting lack of power sources, WSN motes are usually battery powered devices [ML09]. However, batteries impose a finite energy store, which severely limits the achievable performance level and/or system lifetime. This is particularly true for application scenarios with high energy demands, e.g., the operation of a power-hungry GPS receiver [BSB12], since the achievable lifetime is proportional to the ratio of battery capacity and energy consumption. Therefore, to increase the system lifetime, one may equip the mote with a battery of higher capacity. However, increasing the battery capacity may not be feasible for cost, and/or form factor limitations, as a battery's physical size tends to grow proportionally with its capacity. Another approach to extending the lifetime of a system is by lowering its energy consumption. This can be accomplished by reducing the system's power

dissipation and/or duty-cycling, i.e., enable individual system components according to a defined schedule. However, reducing the system's energy footprint by enforcing aggressive duty-cycles results in reduced system availability and performance, which generally conflicts with the application goals, i.e., maximized system utility.

To overcome the lifetime restrictions imposed by purely battery operated systems, and so approach dependable and uninterrupted operation over time periods of multiple years, ambient energy harvesting, particularly in the form of solar energy harvesting has been established [SK11]. While the design of battery operated systems is well understood, the design and operation of energy harvesting systems presents new challenges both in terms of hardware and software. In fact, as will be shown in this thesis, for energy harvesting systems that aim at long-term operation, a paradigm shift in the design and runtime management of the power subsystem is required. Rather than predicting or reacting to the source's short-term variations, the long-term dynamics should be considered both for dimensioning the power subsystem and devising a dynamic energy management scheme that can maximize the usage of available energy [BSBT14b, BSBT14a].

1.2 Energy Harvesting in WSNs

One main advantage of WSNs is that they do not require pre-established infrastructure to function. However, motes are severely power and energy constrained: The instantaneous power dissipation is limited by the battery performance, and the total energy consumption is restricted by the battery capacity. Consequently, the battery presents a major bottleneck that implies a limited achievable lifetime, or, alternatively, requires periodic system maintenance to replace depleted batteries. However, replacing the batteries may not be feasible due to the large number of motes deployed, and/or restricted access to the deployment site. Ambient energy harvesting therefore presents a promising solution to improve (i) system performance, (ii) lifetime, and (iii) maintenance-free autonomous operation. However, as will be discussed in the remainder of this section, simply adding energy harvesting capabilities to the system generally does not suffice to achieve these possible improvements.

In *Section 1.2.1* we give a short overview of energy harvesting techniques. In *Section 1.2.2* we discuss the benefits that can be achieved by leveraging renewable energy sources. In *Section 1.2.3* we point out several challenges that arise in the design and operation of energy harvesting systems, therewith motivating the end-to-end power management solution proposed in this thesis.




Fig. 1.1: (a) Wireless Sensor Network (WSN) mote from Jeelabs powered by a AA Battery [Wip10]. (b) PermaDAQ WSN mote for Permafrost research powered by a D cell battery [BGH⁺09].

1.2.1 Ambient Energy Harvesting

In the context of WSNs we are mainly concerned with ambient energy sources, i.e., solar, wind, vibration, temperature gradients, *etc.* The energy generated by these sources can be converted to electrical energy by use of an appropriate energy transducer. Examples of such energy transducers include photovoltaic cells, wind turbines, piezoelectric devices, and thermoelectric generators. Details on energy conversion is beyond the scope of this thesis, reviews of energy harvesting sources and transducers can be found in e.g., [SK11, CC08].

Naturally, depending on the source, the frequency and magnitude of energy availability can vary greatly. Thus, to sustain a minimum expected system operation, it is crucially important that the power subsystem is provisioned adequately for the respective source, *and* the system has a mechanism to plan and schedule the energy usage. For outdoor WSN applications, photovoltaic energy harvesting with the sun as energy source, i.e., solar harvesting, has been established as the most promising energy harvesting approach. Despite being an uncontrollable source [KHZS07], it follows a well understood diurnal and annual cycle, knowledge of which can be leveraged for designing the power subsystem, and planning the system's energy usage at runtime. Moreover, among the common energy harvesting techniques mentioned, solar harvesting yields the highest power density [CC08]. The mechanically robust energy transducer, which lacks moving parts, further reduces the possibility of failures in the power supply. A failure in the power subsystem can have catastrophic effects in terms of system reliability. Due to the prevalence of solar harvesting in environmental WSN scenarios, this thesis focuses on the design and runtime management of the power subsystem for solar energy harvesting embedded systems.



Fig. 1.2: A solar energy harvesting WSN mote packaged in a weather proof case [RKH⁺05].

1.2.2 Benefits of Solar Energy Harvesting

Over the past decade, energy harvesting has gained significant interest for application in outdoor WSN deployments. There are a number of reasons for this. First and foremost, the prospective of continuous system operation over time periods on the order of multiple years enables previously impossible application scenarios. For example, many researchers in environmental sciences, e.g., Permafrost research [WBB⁺13, WGGP15], rely on monitoring of processes over very long time periods, and therefore benefit from multi-year sensing. Further, aside from the operational advantage of extended lifetime, energy harvesting can significantly reduce the system cost in terms of packaging, development, as well as installation and maintenance. The battery's physical size, which is roughly proportional to the capacity, is often the major contributor to the battery-powered WSN mote's size, see for example *Figure 1.1*. Considering that, for energy harvesting systems, the battery acts as an energy buffer, rather than an energy source, the required battery capacity, and therefore its physical size, may be reduced without impacting the system operation (see *Figure 1.2*). Moreover, the reduction in size not only affects the mote's overall component cost, but also deployment cost, as the individual motes may become significantly smaller and lighter. On the flip-side, energy harvesting incurs additional expenditures in terms of components (e.g., energy transducer), and design and implementation (e.g., control circuitry). Considering the commercial availability of low-cost solar panels and charge control circuitry, the benefits of energy harvesting tend to outweigh the overhead. The challenges in design and operation of solar energy harvesting embedded systems are further discussed in the following section.

In outdoor wireless sensor network deployments, solar energy harvesting with photo-voltaic energy transducers, i.e., solar panel, is most commonly employed. There are a number of reasons for the prevalence of solar energy harvesting. First, solar panels achieve a high energy density, even in indoor scenarios [Ran06]. Low-cost commercial availability and

relatively high conversion efficiency [Jam08], reliability, and a periodically recurring energy source – the sun – are further contributing success factors. The cyclic behavior of the sun guarantees periodic opportunities for recharging the batteries. In theory, this enables perpetual system operation, which eliminates the need for periodic and costly battery replacement trips.

1.2.3 Challenges in Design and Operation of Solar Energy Harvesting Systems

As mentioned in the previous section, energy harvesting systems incur additional cost in terms of harvesting and storage components, as well as control and supervisory circuitry. Further, an efficient utilization of the hardware implies appropriate software control, such that safe battery usage can be achieved while minimizing charging/discharging related battery deterioration. Clearly, the hardware should be designed properly, keeping design trade-offs in mind. For example, while a Maximum Power Point Tracking (MPPT) charge controller can maximize the harvestable energy [BBMT08], it can incur significant cost and operational overhead. This means that the system's entire energy budget must be considered to determine whether the added complexity can be amortized, as suggested in e.g., [BBMT08, CVS⁺07], or if the extra expenditures, both in terms of component cost and energy consumption, outweigh the benefits, as argued in e.g., [RC06, RGS06, TJC08]. Similarly, the software overhead to manage the hardware must be carefully weighed against the benefits in system utility and longevity. While the hardware and software design presents important considerations and trade-off decisions, in this thesis we are not concerned with intricacies of digital system design, but focus on (i) appropriately provisioning the power subsystem (i.e., panel and battery), and (ii) a light-weight approach to dynamically adapting the system performance at runtime.

WSNs are generally viewed as distributed systems, in which the ensemble of motes performs the intended task collectively. Hence, attempts have been made at making the network communication protocols energy harvesting aware, e.g., [ZRLM09, VST13, CSSJ11]. The fundamental idea is to make routing decisions based on the knowledge of energy harvested at the individual motes and so improve the overall network lifetime. In the case of solar powered networks, this can be achieved by ensuring that the motes that are exposed to direct sunlight can off-load the motes with reduced harvesting opportunities due, e.g., shading. Shading can originate from transient weather conditions, i.e., clouds, as well as fixed structures such as buildings, trees, and other obstructions. However, only considering network-wide optimization can potentially lead to the mote

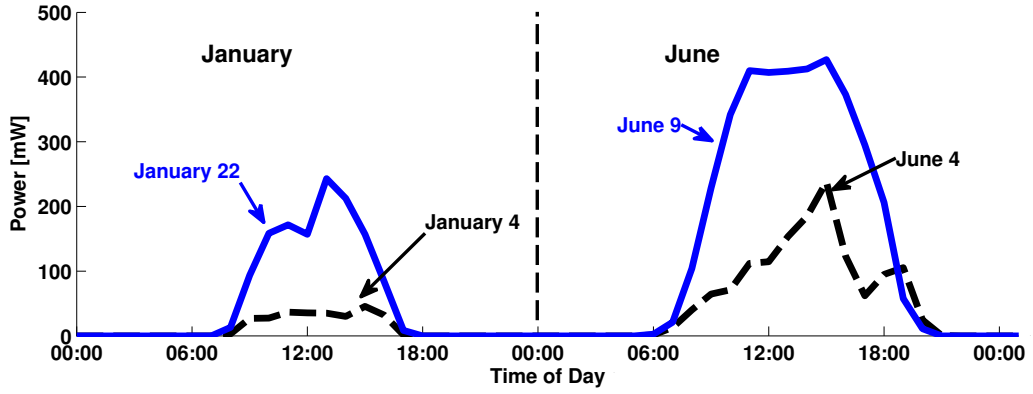




Fig. 1.3: Power profile for 4 days in January and June respectively (CA dataset, see Table 2.1), illustrating significant difference in harvestable energy for different months (due to the earth's orbit around the sun) and within the same month (due to weather).

with the lowest energy availability dictating the entire network's performance. On the other hand, individual motes are often able to perform their intended task even if they are temporarily disconnected from the network by logging the sensed data to non-volatile storage, and communicating the stored data upon re-connecting to the network. Therefore, optimally, the energy usage should be optimized locally, i.e., on a per mote basis, *and* for the entire distributed system [MBTB07]. Nevertheless, in this thesis we are concerned with local optimization by enabling the system to deal with temporal variations of harvestable energy experienced at a single mote, rather than the spatial variations between all the motes within the network.

Figure 1.3 exemplifies the temporal variations of solar energy observed by an outdoor installation of a solar panel in Los Angeles, California (see Table 2.1 for details). By comparing the power profile of the harvested energy on January 4 with that on January 22, it is evident that the source exhibits high short term variations, which originate from varying weather conditions. Similarly, when comparing the energy harvested in January and June, it is obvious that the harvested energy varies seasonally, which is related to the relative position of the earth with respect to the sun. Moreover, as is evident from Figure 1.4, the harvested energy can only be considered constant on average from a long-term perspective (Figure 1.4b). Additionally, the average power observed takes significantly different values depending on the time period considered *and* the instant in time at which the observation period begins (Figure 1.4a). From a short-term perspective, however, the source is highly unstable, as illustrated in Figure 1.3 and Figure 1.4a.



(a) The average accumulated energy observed over 30 days during January differs significantly from that observed over 30 days in July.



(b) The average accumulated energy observed over 10 years exhibits a drastically different profile than when considering the short-term variations shown in Figure 1.4a.

Fig. 1.4: Evolution of accumulated energy harvested over the given time span with a specific harvesting setup (CA dataset, see Table 2.1). A representative indication of average harvestable energy can only be obtained by considering multiple years.

Therefore, in order to deal with the highly non-deterministic short term variability of the source, a feasible power management solution relies on a reasonably accurate prediction of the expected future energy input. For this reason, numerous prediction schemes have been proposed in literature. These schemes, to name a few, include Exponentially Weighted Moving Average filters (EWMA), e.g., [KS03], Weather Conditioned Moving Average (WCMA) [PBAR09], leveraging weather forecasts, e.g., [SGIS10], as well as predictions based on an astronomical model, e.g., [TJC08, BSBT14a]. The dynamic power management solution then computes the future use of energy based on the prediction for a particular time slot. The draw-back of presently proposed solutions lies in the limited prediction accuracy for time horizons exceeding a few minutes, therewith eliminating its application to scenarios where long-term operation is of utmost importance.

In application scenarios where dependable and uninterrupted long-term operation at a deterministic minimum performance level is expected,

a new approach to the design of power management schemes is required. While low-power design principles, and power saving techniques commonly used in battery operated systems are indispensable, we further require a harvesting aware power management that can adapt to the high variability both in time and amount of energy availability.

For example, an approach that greedily uses available energy in an attempt to maximize the utility instead of provisioning for times of deficit may cause system outages at a later time. Not only does this incur high penalties due to low-voltage re-connect hysteresis [BSBT14b], but it also violates the minimum performance objective. On the other hand, conservative short term usage may lead to low minimal service, and the risk of battery overflow. A battery is said to overflow when the battery is full and the surplus energy exceeds the maximum system consumption. This significantly reduces the energy efficiency because the excess energy can neither be stored nor consumed. However, wasting energy is undesirable as it negatively affects the overall system utility.

Therefore, a harvesting-aware dynamic power management scheme aims at scheduling the energy such that the system utility can be efficiently maximized without risking battery over- or underflows, i.e., no energy is wasted due to overflow, and no breakdowns due to underflow. The resulting operating mode is commonly called energy neutral operation [KPS04]. Informally, a system is said to operate in an energy neutral mode if the energy consumed over a given time period is less than or equal to the energy harvested during the same time period. As is intuitively clear, uninterrupted operation over long time periods is only possible when energy neutral operation can be achieved.

In summary, this thesis addresses important questions in the design and operation of solar energy harvesting systems, summarized in the following.

- How should the power subsystem, i.e., solar panel and energy storage element, be provisioned such that the system can optimally deal with the high variations of the underlying source.
- At what point should the system use energy to perform its task, and when should it reduce its energy demand so to conserve energy and/or enable recharging the energy storage element?
- If the energy storage is depleted, the system must shut down and can not perform its intended task. How can such storage underflows be avoided?

- Similarly, if the storage element is full, the harvested energy can not be stored, resulting in wasted energy resources. How can such storage overflows, i.e., wasted energy, be avoided, while ensuring useful utilization of the available energy?
- By definition, WSN motes are highly resource constrained in terms of processing power and memory. Consequently, this thesis addresses implementation details such that the resource constraints can be met while simultaneously improving the overall system utility.

1.3 Review of Power Management Strategies

This section presents a brief overview of State-of-the-Art approaches to power subsystem capacity planning in *Section 1.3.1*, and dynamic power management for energy harvesting systems in *Section 1.3.2*. A review of battery State-of-Charge approximation techniques is deferred to *Section 5.2*.

1.3.1 Capacity Planning

Many design examples of energy harvesting wireless sensing systems can be found in literature, e.g., ZebraNet [ZSLM04], Helimote [RKH⁺05], Ambimax [PC06], Fleck [SCO⁺07], Rivermote [GHSW10], *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 [TJC08], efforts have primarily focused on mitigating the impacts of an inappropriately provisioned power subsystem with energy prediction schemes, e.g., [LW12, SSIS11, BMN09], and dynamic load scheduling based on short-term predictions [PBAR09, VGB07, LSB⁺12, KDMB12]. Three notable exceptions are [KHZS07, TJC08, JC12], which are briefly discussed in the following.

In [KHZS07] 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 Helimotes over two months during Summer in Los Angeles. The approach relies on the availability of a representative energy generation profile and the known system consumption to compute the battery capacity. The limitations of this approach are two-fold. First, the input trace, i.e., energy profile, must

be representative of the conditions at the intended deployment site, and cover at least one full annual solar cycle to yield a suitable battery capacity. Second, the panel size is not considered a design parameter, thus preventing the designer from optimizing the power subsystem with respect to cost, physical form-factor, *etc.* Nevertheless, 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 *Section 2.5* and used as a baseline for evaluation throughout this thesis.

Design experiences of the HydroSolar micro-solar power subsystem are presented in [TJC08]. Despite leveraging the same astronomical model [DHM75] as used in this thesis, 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. 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 work that is closely related to this thesis in that it relies on a similar astronomical model is presented in [JC12]. 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 [TJC08] are used, which may result in an underprovisioned power subsystem.

Adequate capacity planning is commonly accepted as an important design criteria. Nevertheless, little effort has been expended to devise a systematic approach that enables a designer to determine the appropriate battery capacity and panel size for a given application scenario. The only approach that we are aware of [KHZS07], relies on the availability of representative energy traces, which are usually not available. In an attempt to solve this problem, we propose a systematic approach to power subsystem capacity planning for solar energy harvesting embedded systems, such that uninterrupted operation over multiple years at a predefined performance level may be achieved [BSBT14b]. This capacity planning algorithm, which is fully described in *Chapter 2*, leverages a modified astronomical model to approximate the harvestable energy, and thus requires no representative energy traces or other calibration data. The algorithm takes as input the deployment site's latitude, the panel orientation and inclination angles to compute the appropriate battery capacity. An

indication of the expected meteorological and environmental conditions at the intended deployment site is the only unknown input parameter. While this parameter can be easily obtained if energy traces exist, we discuss how it can be obtained when traces are not readily available. The modified astronomical model's ability to predict the harvestable energy is validated 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 capacity-planning algorithm enables uninterrupted system operation, i.e., 100% availability, at up to 53% smaller batteries when compared to the State-of-the-Art.

1.3.2 Dynamic Power Management

In the seminal work on energy harvesting theory [KS03, KPS04, KHZS07], the first dynamic duty-cycling scheme for solar energy harvesting systems was proposed within a theoretical framework that defines *Energy Neutral Operation* as the fundamental requirement of energy harvesting systems so to sustain perpetual operation. *Energy Neutral Operation* is achieved if the system never consumes more energy than what it can harvest over a given time period δ , i.e., the battery fill-level $B_{fill}(t + \delta)$ is greater than or equal to $B_{fill}(t)$. With their approach, a day is discretized into slots of equal duration δ , and the expected energy input for each slot is learned with an Exponentially Weighted Moving Average (EWMA) filter. Each slot's respective duty-cycle is then computed by considering the mismatch between expected and actual energy input. EWMA filters have long been used [Hun86] for data prediction because of its simplicity and relatively good accuracy. However, due to limited correlation between past and future weather conditions, this approach achieves acceptable prediction accuracy only for prediction windows on the order of hours.

Weather Conditioned Moving Average (WCMA), proposed in [PBAR09], improves upon EWMA's prediction accuracy. The authors not only consider the harvested energy in the same time slot during previous days, but also incorporate current weather conditions to obtain the expected energy input in the current slot. While achieving an almost three-fold improvement in prediction accuracy over EWMA, it is not clear if and how this improvement translates into increased system performance and/or energy neutrality. Just as the EWMA approach, WCMA is also constrained by short prediction windows, and thus tends to suffer from high duty-cycle variance.

More recently, the use of professional weather forecast services have been considered to predict the disposable energy [SGIS10]. The authors formulate a model to translate weather forecasts into solar or wind energy harvesting predictions. While it is unclear what baseline is used, the

authors conclude that their energy predictions are more accurate than those based on past local observations.

In [LSB⁺12] and [VGB07], model-free approaches to dynamic performance scaling are presented. In [VGB07], a technique from adaptive control theory, i.e., Linear-Quadratic Tracking, is used to dynamically adapt the system's duty-cycle based on the battery State-of-Charge and so ensure *Energy Neutral Operation*. For the datasets evaluated, the authors report between 6 and 32% improvement in mean duty-cycle, and between 6 and 69% reduction in duty-cycle variance when compared to EWMA. Similarly, in [LSB⁺12] a Proportional-Integral-Derivative (PID) controller monitors the energy storage element, and dynamically adapts the duty-cycle such that an expected voltage level of the storage element (a super-capacitor in this case) can be maintained. While presenting low-complexity solutions, both of these approaches suffer from high duty-cycle variability, and rely on a well performing battery State-of-Charge approximation algorithm. The PID approach additionally requires parameter tuning, for which, however, solutions exist in the literature.

In contrast to the above solutions, which tend to focus on dealing with the short-term energy variations, in *Chapter 3* we present a new and promising approach that takes uninterrupted long-term operation at a predefined minimum performance level as the ultimate goal. This scheme takes the astronomical model discussed in *Chapter 2* to provision the power subsystem, i.e., battery and panel, and leverages the same model at runtime to adjust the utilization such that long-term energy neutral operation may be maintained. Through extensive simulation we show that the approach can adjust to deviations from the design time expectations, while maintaining a significantly higher minimum service-level than the approaches discussed above.

The problem of optimal energy allocation, which we address in *Chapter 4*, is closely related to the work in [CSSJ11]. The authors consider the problem of maximizing the throughput of a network of sensor nodes powered by harvesting energy. They study the problem for a single node, and then as a distributed algorithm for multiple nodes. They establish a relation to the shortest path problem in a simply connected space for a single node with known harvested energy. For multiple nodes they propose a heuristic that is optimal under the assumption that each node receives *homogeneous* harvested energy. In contrast to [CSSJ11], in *Chapter 4* we study the objective of maximizing the minimum used energy over all time steps for a single node. Interestingly, we also establish a relation between this problem and a Euclidean path minimization problem. Further, we study how to model and factor the expected variability in the input harvested energy. In contrast to [CSSJ11], we also extensively validate our approach with real-world data from known databases.




Fig. 1.5: Chapter structure of this thesis. The energy availability model developed in *Chapter 2* will be used in *Chapters 3* and *4*. The dynamic power management schemes developed in *Chapters 3* and *4* as well as the battery State-of-Charge approximation algorithm will be further evaluated using a real-world WSN deployment.

1.4 Thesis Outline and Contributions

In this thesis we present a novel end-to-end power management solution for solar energy harvesting systems consisting of (i) power subsystem capacity planning, and (ii) dynamic power management. We will use solar energy harvesting Wireless Sensor Networks (WSNs) as exemplary application scenario, but note that our approach is applicable to other application scenarios that leverage a periodic energy source. *Figure 1.5* presents an overview of the chapter structure of this thesis. We will start with a discussion of power subsystem capacity planning in *Chapter 2*, and present a systematic approach such that an appropriate solar panel size and battery capacity for a given application scenario can be obtained without the need for solar trace data. Then, assuming a suitable power subsystem has been defined, in *Chapter 3*, we investigate a runtime algorithm that aims at ensuring dependable and uninterrupted long-term system operation at a predefined minimum performance level. In *Chapter 4*, we present the first formal study on optimizing the energy utilization of energy harvesting embedded system while giving bounds on the minimum energy usage. We further define an optimal energy scheduling algorithm that can give guarantees on minimum energy utilization. In order for the dynamic power management scheme to be able to make informed decisions on future energy usage, the system requires a means to obtain the residual charge stored in the battery. Hence, in *Chapter 5* we present a light-weight battery State-of-Charge approximation algorithm that can provide accurate indications of residual charge without requiring hardware support that exceeds standard system monitoring circuitry. Finally, in *Chapter 6*, we further illustrate the benefits of our end-to-end power management approach using a real-world sensor network deployment in the Swiss alps [BBF⁺11]. Throughout this thesis we assume a




Fig. 1.6: System architecture assumed throughout this thesis. The system is partitioned into harvesting subsystem (solar panel), storage subsystem (battery and charge controller) and the load to be supported (wireless sensing system). The flow of generated energy is indicated with dotted lines, while the dashed lines represent the energy consumed.

system architecture as illustrated in *Figure 1.6*, where we combine the *Harvesting* and *Storage* subsystems into the power subsystem. The *load* is assumed to be a WSN mote that executes a given task, which results in an energy consumption that is characteristic of the task. The system model will be refined as required in the following chapters.

A brief introduction, and summary of contributions of the individual chapters is presented in the following.

Chapter 2: Power Subsystem Design

Chapter 2 presents and evaluates a novel capacity planning algorithm to compute the battery capacity and solar panel size necessary for a given application scenario and expected geographical deployment location. In contrast to the State-of-the-Art, no energy traces are necessary for the algorithm to obtain an adequate power subsystem specification. This is very desirable, as representative energy traces for a particular location may not be readily available. The technique discussed in this chapter can be used as design space exploration tool that enables the system designer to investigate size, weight and power trade-offs and so identify the most cost-efficient combination of battery and solar panel for a particular application. The energy availability model developed in this chapter is later exploited in *Chapters 3* and *4* for devising dynamic power management schemes that enable dependable and uninterrupted long-term operation for state-of-the-art energy harvesting embedded systems. In summary, *Chapter 2* makes the following contributions.

- We propose a novel power subsystem capacity planning algorithm based on a modified astronomical model [DHM75] that serves as the energy availability model. The benefit of using the astronomical model lies in the fact that it eliminates the need for detailed power traces or calibration data.
- We validate the energy availability model's ability to approximate the harvestable energy with power measurements of a solar panel, and show that it performs very well at predicting long-term energy availability.
- Through extensive simulation we evaluate the capacity planning approach, and show that it yields smaller batteries than the State-of-the-Art [KHZS07] capacity planning algorithm without impacting minimum system performance.
- We further show that our approach outperforms the State-of-the-Art by achieving 100% system availability over ten years for three different data traces, while requiring up to 53% smaller batteries for the evaluated datasets.

Chapter 3: Long-Term Dynamic Power Management

While it will be shown that the capacity planning algorithm introduced in *Chapter 2* computes a suitable power subsystem specification, achieving long-term operation is only possible if the design-time assumptions do not deviate significantly from the true energy conditions observed at the intended deployment site. Therefore, in *Chapter 3* we present and evaluate a novel algorithm to dynamically adjust the system duty-cycle of solar energy harvesting embedded systems such that uninterrupted operation, even in the presence of large deviations from design-time assumptions, can be sustained over time periods on the order of multiple years. Contrary to previously proposed techniques, our approach takes (i) a long-term view of the energy source, and (ii) explicit knowledge of the power subsystem capacity to satisfy energy neutrality. In this chapter we leverage the results from *Chapter 2* to devise a dynamic power management scheme that maximizes both minimum and average duty-cycle, as well as duty-cycle stability. More specifically, the contributions of this chapter is given in the following.

- We show how the energy availability model from *Chapter 2* can be adapted such that it serves as a suitable online energy predictor with little runtime overhead.

- We present a new dynamic power management algorithm that uses the online energy predictor to compute the long-term sustainable performance level, i.e., duty-cycle, at runtime.
- We show that dynamic power management significantly improves system utility when compared to applying appropriate design-time capacity planning only.
- We show that the duty-cycle computed with our approach exhibits very low variability, despite relying on, and reacting to a highly variable energy source. Achieving high duty-cycle stability can be a strong requirement in a broad range of application scenarios, e.g., [VKR⁺05, CB10, WBB⁺13].
- Through simulation with eleven years of data at three different geographical locations we show that our algorithm outperforms the State-of-the-Art in energy-predictive [KHZS07, PBAR09], and battery-reactive [VGB07] performance scaling approaches in all relevant performance metrics.
- Extensive performance evaluation shows that our approach improves on the State-of-the-Art approaches in terms of average sustainable performance level by up to 177%, energy efficiency by up to 184%, and duty-cycle stability by up to three orders of magnitude, while incurring zero downtime, i.e., system availability of 100%.

In order to investigate the approach from this chapter in a realistic setting, and to validate the simulation framework, we present a case study in *Chapter 6*. The case study evaluates the performance of our approach using an *X-Sense* environmental monitoring system [BBF⁺11] deployed over two years in a remote and inaccessible high-alpine environment. We further demonstrate that significant improvements in system utility can be achieved without risking downtime due to power outages.

Chapter 4: Optimal Power Management

In *Chapter 4* we present the first formal study on optimizing the energy utilization of energy harvesting embedded systems while simultaneously providing bounds on the minimum achievable energy utilization. To this end, we formally study the energy harvesting problem with the objective of maximizing the minimum energy used across all time intervals. We first consider a simplified clairvoyant setting wherein the harvested energy for a given time interval is exactly known in advance. With this assumption, we derive and prove an optimal algorithm that computes the amount of energy to be used in each time interval, such that the minimum energy used is maximized over all intervals. To identify an efficient

algorithm to compute the optimal use function, we establish a relation between the energy harvesting problem and the shortest Euclidean path problem. This relation allows us to use well-studied and efficient algorithms for the clairvoyant energy harvesting problem. Then we study the more general and realistic problem in which the harvested energy is not known exactly; instead we know only a conservative estimate of it. With this assumption, we show how a finite horizon scheme can be used to adaptively update the energy utilization at runtime. We prove that, under certain realistic assumptions of the energy estimate, such a finite horizon scheme is guaranteed to provide a certain minimum energy usage that is better than a non-adaptive scheme.

The contributions of *Chapter 4* can be summarized as follows.

- We define optimality criteria that an optimal power management controller must satisfy.
- We present and prove the optimality of a clairvoyant algorithm to compute the optimal use function such that (i) the minimum utilization is maximized, and (ii) the total system utility is maximized.
- We adapt the clairvoyant algorithm using a finite horizon scheme to the realistic scenario in which the harvested energy is not known a priori.
- We prove that the finite horizon scheme achieves a guaranteed minimum energy utilization that is larger than what can be achieved with a non-adaptive scheme.
- We show that the energy availability model from *Chapter 2* can be adapted with a non-uniform linear scaling function, and show that a single scaling function can be used to obtain a suitable energy predictor for 8 different geographical locations.
- We describe a look-up table based implementation of the finite horizon scheme, which results in very low runtime overhead.
- With extensive experimental evaluation using 8 publicly available datasets, we quantitatively establish that the proposed solution significantly outperforms four previously proposed State-of-the-Art solutions. In fact, at best, our approach reaches to within 9.9% of the theoretical optimal performance.

Just as with the algorithm from *Chapter 3*, the power management algorithm presented in this chapter will be further evaluated using the case study in *Chapter 6*.

Chapter 5: State-of-Charge Approximation

In *Chapter 5* we present a light-weight and cost-effective approach to approximating the battery State-of-Charge based on voltage measurements only, which is highly desirable for systems with limited hardware support. It has been shown that awareness of the energy available to the individual motes [PLR05], and the entire network [CD05] can significantly improve overall system lifetime and utility. However, accurately determining the battery fill level, referred to as battery State-of-Charge, presents a non-trivial problem to solve. This is because a battery's State-of-Charge depends on many battery internal and external factors, such as size and type of battery, the rate at which it is discharged, as well as temperature and battery condition (e.g., age, present State-of-Charge, *etc.*). While knowledge of the State-of-Charge may not be mission-critical for traditional battery powered WSN applications, improved observability, predictability, and utility [PAG09, TJC08] tend to outweigh the efforts in implementing the necessary functionality for energy harvesting systems. In fact, accurate approximation of the battery State-of-Charge is crucial for harvesting-aware dynamic power management schemes such that a safe, but competitive system duty-cycle can be computed. We show that with our approach, State-of-Charge approximations with up to 95% accuracy are possible, while taking battery inefficiencies due to e.g., temperature and aging into consideration, despite not explicitly modeling these effects.

In summary, *Chapter 5* presents the following contributions.

- We present a practical, trace-based, direct-measurement [BKN02] method for online battery State-of-Charge approximation that does not rely on special purpose hardware.
- We show how leveraging known characteristics of the solar panel and battery behavior under load can be used to devise a computationally light-weight method to State-of-Charge approximation.
- We present a state machine for tracking the charging and discharging phases at runtime.
- We show how the battery can be profiled such that the model parameters can be defined even for batteries that differ from the profiled battery in terms of battery age and capacity, as well as operating temperature.
- We perform extensive experiments under different conditions (operating temperature, load variations, battery age).

- Using the experiments, we show that our direct-measurement approach achieves State-of-Charge approximation with an average error below 5% when compared to discharge tests.
- Finally, we show that accurate lifetime predictions even with temperature fluctuations and varying battery and load conditions are possible, as our method implicitly accounts for battery inefficiencies due to e.g., temperature and aging.

Chapter 6: Case Study

In *Chapter 6*, we consider a real-world environmental monitoring scenario that requires uninterrupted system operation over time periods on the order of multiple years. To achieve this goal, we leverage the techniques discussed in the previous chapters. More specifically, we demonstrate the benefits of the proposed end-to-end power management solution, i.e., power subsystem capacity planning (*Chapter 2*), dynamic power management schemes (*Chapters 3 and 4*), and the State-of-Charge algorithm (*Chapter 5*), using the *X-Sense* wireless sensor systems [BYL⁺11, BSB12] deployed for on-site design-space exploration in a remote, high-alpine environment as a case study [BBF⁺11]. A performance evaluation over two years reveals that dynamic power management significantly improves the system utility when compared to only applying appropriate capacity planning.

The key results obtained with the case study are summarized as follows.

- We show how the battery State-of-Charge algorithm from *Chapter 5* can be adapted to provide the information required by the dynamic power management algorithms.
- We describe a minor modification to the algorithm from *Chapter 3* such that the surplus energy during summer can be better utilized. This is necessary as the *X-Sense* platform's power subsystem is inadequately provisioned.
- We show that with the real-world datasets, the optimal algorithm from *Chapter 4* reaches the minimum utilization to within 21.1% of the theoretical optimal, and the average utility to within 8.3% of the optimal.
- Finally, we show that dynamic power management yields significant improvements over only applying capacity planning, while exhibiting robustness against variations in the observed energy profile, irrespective of the source of deviations from the model.

2

Power Subsystem Capacity Planning

2.1 Introduction

Wireless Sensor Networks (WSNs) are predominantly deployed in remote and inaccessible sites. Due to lack of power sources at these sites, the motes comprising such a network are usually battery powered devices. However, the finite energy store imposed by non-rechargeable batteries severely limits the lifetime and achievable performance level of application scenarios with increased energy demands, e.g., [BSB12]. Ambient energy harvesting, particularly in the form of solar energy harvesting [CC08, SK11], has thus attracted much attention as a promising solution for enabling perpetual system operation.

However, simply enhancing a system with energy harvesting capabilities may not suffice to achieve uninterrupted long-term operation at a predefined minimum performance level [HG10, STC13]. This is because solar energy harvesting opportunities depend both on static and dynamic factors, e.g., efficiencies of the solar panel, energy storage element and charge controller, solar panel installation parameters [HG10], 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 size 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 chapter addresses this issue and provides a tool

for systematic design space exploration and identification of the smallest combination of battery and solar panel for a particular application.

Challenges

While the design of the power subsystem for battery powered electronic systems has been relatively well understood, appropriately provisioning the battery capacity and solar panel size for a given application scenario and environmental conditions presents new challenges. Ultimately, the performance level achievable by an embedded computing system is 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 can be difficult to define either of these periods without resorting to overly pessimistic assumptions [TJC08].

In fact, literature review shows that most efforts attempt to mitigate the effects of an inadequately provisioned power subsystems with dynamic load scaling schemes, e.g., [VST13, LSB⁺12], which are based on expected harvesting opportunities predicted by an energy prediction scheme, e.g., [CPS12, VGB07, PBAR09]. 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 power subsystem capacity planning. 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, runtime approaches alone are unsatisfactory because they do not consider the limitations of the power subsystem.

Contributions

To alleviate the above 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 information, i.e., geographical location and panel orientation and inclination angles. 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., shadowing due to growing vegetation, dirt, *etc.*. For scenarios where the observed conditions at runtime deviate significantly from expected conditions, the dynamic power management techniques discussed in *Chapters 3 and 4* ensure that the system can react to the variations and so maintain the minimum performance objective.

The contributions of this chapter are summarized as follows.

- Firstly, we propose a power subsystem capacity planning algorithm based on an astronomical model [DHM75] that we adapt to obtain an energy availability model that can approximate the harvestable energy for a given geographical location and harvesting setup.
- Secondly, we perform measurements with a solar panel to verify the energy availability model's ability to approximate the harvestable energy, and show that it is sufficiently accurate for estimation of long-term solar energy harvesting opportunities.
- Thirdly, we evaluate the capacity planning approach through simulation, and show that it yields smaller batteries than the State-of-the-Art capacity planning algorithm [KHZS07].
- We further show that our approach outperforms the State-of-the-Art by achieving 100% availability over ten years for three different test datasets, while requiring up to 53% smaller batteries. In contrast to the State-of-the-Art, 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.

Roadmap

The remainder of this chapter is structured as follows. In *Section 2.2* we introduce the high-level concept of the power subsystem capacity planning approach. Then in *Section 2.3* we provide a detailed discussion of the harvesting conditioned energy availability model that forms the basis of the capacity planning algorithm, which is detailed in *Section 2.4*. In *Section 2.5*, we evaluate the proposed method's ability to achieve uninterrupted long-term operation, and its sensitivity to parameter variations. Finally, *Section 2.7* concludes this chapter with a short summary.




Fig. 2.1: 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. Note, numbers in parentheses refer to sections.

2.2 Concept

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. *Figure 2.1* shows the high-level design flow of the proposed capacity planning approach that satisfies this objective. The dashed boxes represent user inputs that characterize the deployment setup, hardware technology employed, electrical load, and expected meteorological conditions. The capacity planning algorithm discussed in *Section 2.4* then uses these parameters to compute the battery capacity required that can indefinitely sustain 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 [DHM75]. This model, which is fully described in *Section 2.3.1*, 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 *Section 2.3.1* 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*, which are described in *Section 2.3.2*. 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 (η_{in} and η_{out}) of the chosen storage element. Note that in this work we are not concerned with the selection of optimal storage technology and energy conversion architecture, as this is highly application specific, but rather aim at finding the minimum feasible power subsystem as characterized by the above parameters. Further note that we intentionally leave the selection of panel size to the designer and compute the battery capacity for the given panel and application requirements. This is not a limitation, but rather a choice, as our approach can easily be adapted to compute the panel size required for a given battery capacity.

Finally, the *system parameters*, i.e. power dissipation P_{sys} , and minimum expected duty-cycle DC_{sys} , as well as power dissipation by power-conditioning circuitry, i.e., power dissipation P_{vr} and P_{cc} of the voltage regulator and charge controller respectively, characterize the load imposed on the battery. The system's total energy requirement then 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} . Ideally, the computed duty-cycle $DC_{supported}$ is equal to expected duty-cycle DC_{sys} . However, when discrete panel sizes are considered, the supported duty-cycle may be slightly higher or lower than the duty-cycle DC_{sys} desired by the designer. Note that P_{sys} and DC_{sys} may be a vector of length M , specifying the power dissipation and duty-cycle of the M system components. Further note that in this chapter we assume constant power dissipation by the load (but not necessarily constant energy consumption); 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 *Section 2.3.1* can accurately approximate the long-term energy harvesting opportunities.

2.3 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 [DHM75]. *Section 2.3.1* discusses three modifications to this model such that varying environmental conditions can be taken into consideration. *Section 2.3.2* describes the load

model and the harvesting conditioned energy availability model, which incorporates conversion and storage inefficiencies.

2.3.1 Energy Availability Model

According to [DHM75], the total solar energy $E_{astro}(\cdot)$, incident on a flat surface is a function of the time of the year, the area of the surface, the deployment location's latitude L , and azimuth (orientation) and inclination angles ϕ_p , and θ_p respectively. As shown in Equation (2.1), 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 [DHM75]. 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 [DHM75], and [BC06, HLG06] respectively.

$$E_{astro}(\cdot) = E_{sun}(d, t, L, \phi_p, \theta_p, \tau) + E_{sky}(d, t, L, k, \theta_p, \tau) + E_{gnd}(d, t, L, k, R, \theta_p, \tau) \quad (2.1)$$

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 [HLG06]. The authors in [DHM75] 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 [DHM75, BC06]. 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 Section 2.5.1) can be closely approximated with $E_{astro}(\cdot)$ by letting $k = 0.1$ and varying τ , such that $\sum^T E_{astro}(\cdot) \cong \sum^T E_{actual}(d)$, where $T \gg 1$ day. We leverage this observation and define the so called environmental parameter Ω in Equation (2.2) to replace τ . The parameter

δ_i represents the proportion of time during which the atmosphere exhibits the optical thickness τ_i .

$$\Omega = \sum_i^N \delta_i \tau_i, \text{ where } \sum_i^N \delta_i = 1 \quad (2.2)$$

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 *Sections 2.5.2 and 2.6*).

As is evident from *Equation (2.1)*, 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 [GGP12], 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 [DHM75] 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., $R = 0.3$ [DHM75]. 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 $R = 0$.

2.3.2 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 [HG10]. This section discusses the effects of non-ideal harvesting and storage elements.

System Architecture

In this work we assume a harvest-store-use system architecture as defined in [SK11]. 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. However, 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 actually converted to electrical energy. Depending on technology, the panel's conversion efficiency η_{pv} achieves a few percent for thin-film technologies, and exceeds 40% for high-end multi-junction cells [GEH⁺12]. Furthermore, a solar panel has a manufacturer specified maximum output power 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 $E_{pv}(\cdot)$ 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 [WLLP01, HG10, STC13]. Large-scale photo-voltaic installations are usually only used for WSN base stations and experimental units when mains power is not available. Therefore, 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 Equation (2.3), where η_{cc} represent the charge controller's efficiency.

$$E_{pv}(d, \Omega) = A_{pv} \cdot \eta_{cc} \cdot \eta_{pv} \sum_{t=0hr}^{t=23hr} \min \{ 1hr \cdot P_{pv}, E_{astro}(d, t, L, k, R, \theta_p, \phi_p, \Omega) \} \quad (2.3)$$

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, as indicated by the multiplicative factor η_{cc} in Equation (2.3). 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 [TJC08]. The choice of technology is very application specific, and there are arguments advocating advanced charge controllers [BBMT08, CVS⁺07], while others argue that, for micro-solar energy harvesting systems, the gain is dwarfed by the energy expenditure of the controller [TJC08]. It is also possible to operate without a charge controller [PC06], but the lack of over-voltage protection may significantly reduce battery life [CVS⁺07].

Charge controllers often implement a battery protection mechanism, known as low-voltage disconnect [BAB13]. This means that 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} . Details on how we take this feature into consideration are discussed in *Section 2.5.1*. Clearly, while a battery protection mechanism is desirable, prolonged downtime due to protection against deep discharge cycles can incur significant performance penalties. Moreover, deep discharge 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 energy 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, fill-level, and charging strategy [BAB13].

Hence, to account for charging and discharging inefficiencies of the battery during simulation (see *Section 2.5*), the energy flowing into and out of the battery is scaled by the respective efficiency factors η_{in} and η_{out} , as shown in *Equations (2.4)* and *(2.5)*. The loss in the charging process due to battery internal resistance and electrochemical processes is represented by η_{in} . The factor $1/\eta_{out}$ accounts for the fact that only a fraction of the charge transferred into the battery during charging can be recovered when discharging the battery [BAB13].

$$E_{in}(d) = \min \left[\eta_{in} \cdot E_{pv}(d, \Omega), \min \{0, \eta_{out} \cdot B_{nom} - B(d-1) - E_{out}(d)\} \right] \quad (2.4)$$

$$E_{out}(d) = E_{load}(d)/\eta_{out} + E_{leak} \quad (2.5)$$

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 *Equation (2.4)*. Similarly, to support the energy consumption $E_{load}(d)$,

which represents all energy consumers regardless of function (see *Equation (2.7)*), the current battery level and harvested energy must exceed the amount to be withdrawn from the battery, i.e., $B(d-1) + E_{in}(d) > E_{out}(d)$. The battery specific leakage E_{leak} is assumed to be constant [KHZS07]. Then, with the battery characterized by *Equations (2.4) and (2.5)*, the battery State-of-Charge at the end of a given day is obtained with *Equation (2.6)*. B_{nom} is the manufacturer rated nominal capacity, which is generally given in Ampere-hours, but converted to Watt-hours for our purpose. 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., $B(d) \geq B_{min} \forall d$.

$$B(d) = \min \{ \eta_{out} \cdot B_{nom}, B(d-1) + E_{in}(d) - E_{out}(d) \} \quad (2.6)$$

We ignore aging effects of the battery [BAB13]. However, we note that our approach results in very shallow discharge cycles and so protects the expected battery lifetime [BAB13]. In fact, our method results in one full discharge cycle per year; hence, assuming a battery rated for a few hundred discharge cycles [BKN02], the battery is expected to outlast other system components.

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_{sys_i}) and respective active (P_{on_i}) and inactive (P_{off_i}) power dissipation. The total daily energy required to operate at the expected performance level is then defined by *Equation (2.7)*, where $\gamma = 24$ hours.

$$E_{load}(d) = \gamma \cdot \left[P_{cc} + P_{vr} + \sum_i^M (DC_{sys_i} \cdot P_{on_i} + (1 - DC_{sys_i}) \cdot P_{off_i}) \right] \quad (2.7)$$

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.




Fig. 2.2: Actual ($E_{real}(d)$), and modeled daily energy input ($E_{in}(d)$) and output ($E_{out}(d)$) for panel area $A_{pv} = 0.02m^2$ and the CA dataset (see Section 2.5.1). Surplus energy is indicated with the hatched area. The cross-hatched area between d_1 and d_2 shows the required battery capacity.

2.4 Capacity Planning Algorithm

As discussed in the previous section, 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 Equation (2.2), or, if available, by profiling a representative dataset. The total energy consumption, defined in Equation (2.5), completes the necessary information required 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 (see Figure 2.2). 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. One such cycle is illustrated in Figure 2.2, which shows the actual energy input $E_{real}(d)$ observed Los Angeles, California (CA dataset, see Section 2.5 and Table 2.1), the daily system consumption $E_{out}(d)$ to be supported, and the energy input approximation $E_{in}(d)$ such that $\sum^T E_{real}(d) \cong \sum^T E_{in}(d)$ over $T = 365$ days. Note that, depending on the panel installation parameters, i.e., orientation and inclination angles, the harvesting profile may exhibit two peaks per year (see Figure 2.6b) as opposed to just one like in Figure 2.2. This is not a problem, as our algorithm will still ensure that the period of surplus preceding the period of deficit can ensure that the battery is fully recharged before the start of the period of deficit.

The battery capacity, B , required to operate the system during times of deficit, i.e., $E_{out}(d) > E_{in}(d)$, $d \in [d_1, \dots, d_2]$, is given in Equation (2.8) and il-

illustrated by the cross-hatched area in *Figure 2.2*. Note that we assume the battery to be fully charged on day d_1 . The first term on the left-hand side in *Equation (2.8)* specifies the amount of energy that is necessary 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} . Note that inefficiencies are already considered in $E_{out}(d)$ and $E_{in}(d)$ (see *Equations (2.4)* and *(2.5)*).

$$\sum_{d_1}^{d_2} (E_{out}(d) - E_{in}(d)) = B_{nom} \quad (2.8)$$

Similarly, *Equation (2.9)* 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 *Figure 2.2*). As mentioned previously, the harvested energy during periods of surplus must be able to recharge the battery. Therefore, to achieve perpetual operation over multiple years, the left-hand term in *Equation (2.9)* must be at least as large as the left-hand term in *Equation (2.8)*. Note that satisfying *Equation (2.8)* alone does not allow long-term continuous operation, as $E_{out}(d)$ may not be sustained.

$$\sum_{d_0}^{d_1} (E_{in}(d) - E_{out}(d)) \geq B_{nom} \quad (2.9)$$

Then, assuming the panel size is fixed, and we wish to find the minimum required battery capacity B_{nom} , we vary the energy consumption E_{out} (via the duty-cycle) until *Equations (2.8)* and *(2.9)* hold. For example, with the dataset shown in *Figure 2.2*, and a panel size of 20cm^2 and $E_{load} = 227.8\text{mW}$, the required battery capacity is approximately 68Wh (see *Section 2.5.3*). 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 example clearly illustrates the non-linear relationship between battery capacity and solar panel size.

Note that with this approach we can also determine the required panel size when the battery capacity is given. In this case, we vary the panel size, and so change the energy generation $E_{in}(d)$, until *Equations (2.8)* and *(2.9)* hold.

2.5 Evaluation

This section evaluates the proposed method's ability to compute a battery capacity, given a solar panel size, such that uninterrupted long-term operation of solar energy harvesting systems can be ensured.

Tab. 2.1: 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

2.5.1 Evaluation Methodology, Validation Data, and Performance Metrics

Methodology

To validate the energy availability model, we first compare its energy estimation accuracy to measurements performed with a solar panel. As a second step, we perform extensive simulations to evaluate the proposed method's ability to yield a battery capacity that can support long-term operation. The simulation framework simulates a system according to *Sections 2.3* and *2.4* and the input data discussed in the following.

Model Validation Input Data

For validation of the energy availability model, we obtain ground-truth data by measuring the power generated by a $0.1725m^2$ mono-crystalline solar panel (cleversolar CS-30 [cle]) 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 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 2.1*). We use 11 years of data, from which the first year (i.e., days 1-365) of each location is used as calibration

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

data (see *Section 2.5.3*), while the data for the remaining 10 years is used as input for the simulation discussed in *Section 2.5.3*.

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 *Section 2.3.2*. For the battery technology parameters we assume $\eta_{in} = 0.9$, $\eta_{out} = 0.7$, which is commonly done in literature, e.g., [KPS04, TJC08, PBAR09, VGB07]. For the charge controlling and power conditioning circuitry we assume zero loss, i.e., $\eta_{cc} = \eta_{vr} = 1$. Finally, for the panel parameters we assume a typical midrange efficiency of $\eta_{pv} = 10\%$ [GEH⁺12], and zero inclination angle. The reconnect hysteresis (see *Section 2.3.2*) is set at 30% of the battery capacity $B = \eta_{out} \cdot B_{nom}$.

Performance Metrics

For performance comparison between the approach discussed in this chapter and the State-of-the-Art [KHZS07], 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.

2.5.2 Energy Availability Model Validation

The measurement data described in *Section 2.5.1* is used as ground-truth for validation of the energy availability model. Here we are concerned with how well the actual energy input can be approximated with a given environmental parameter Ω . We assume $\Omega = 0.51$ as per *Equation (2.3)* 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 panel's specification, and the technology and deployment parameters from *Section 2.5.1*.




Fig. 2.3: Ratio of total daily energy predicted with our approach ($E_{pv}(d, \Omega = 0.51)$), and actual energy input ($E_{real}(d)$) for each day. Note: bars are capped at 2, and labeled with actual value.

As is evident from *Figure 2.3*, which shows the ratio of estimated and actual energy input for each day, the model tends to significantly overestimate available energy for days with little, to no energy harvesting opportunities. However, for the majority of days the ratio is close to unity, illustrating the model's ability to closely approximate the harvestable energy. *Table 2.2a* lists statistics for the same 41-day period. *Table 2.2b* shows the results when considering an estimation granularity of one week. On average 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 [GGP12].

2.5.3 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 the current State-of-the-Art [KHZS07], we are not aware of any other concrete algorithms for power subsystem capacity planning of energy harvesting systems. Thus, the method described in [KPS04, KHSR06, KHZS07] is used as a baseline and briefly reviewed in the following.

2.5.3.1 Reference Model

In [KPS04, KHSR06, KHZS07] the authors present a set of abstractions for capacity planning of energy harvesting systems that can be considered the

Tab. 2.2: Energy availability model validation statistics.(a) Statistics for E_{pv}/E_{real} for the 41 day experiment shown in Figure 2.3.

Min	Mode	Max	St. Dev.	Mean	Range	Median
0.7947	0.7947	26.87	5.737	3.389	26.07	1.01

(b) Weekly energy sums. E_{real} refers to measurements with solar panel.

	Week						Total
	1	2	3	4	5	6	
E_{pv} [Wh]	4844	5646	5034	5109	5164	3706	29503
E_{real} [Wh]	3773	5775	3155	4929	3605	3255	24492
E_{pv}/E_{real}	1.28	0.98	1.6	1.04	1.43	1.14	1.20

State-of-the-Art in harvesting theory. The authors define Energy-Neutral Operation as a performance metric, and formally state the conditions that must be met to achieve Energy-Neutral Operation. The authors show 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 [KPS04]. Despite defining a battery capacity for long-term continuous operation, the authors conclude that, for achieving Energy-Neutral Operation, 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; instead the system must rely on a well performing energy prediction scheme to achieve acceptable long-term performance.

The main limitation of this approach is its dependence on appropriate energy traces. In order to extract ρ_s , this algorithm requires a dataset that is representative of the conditions at the deployment site. However, when applying the method described in [KPS04, KHSR06, KHZS07] to obtain this quantity, it is found that their approach yields significantly different performance levels depending on the particular time period used (see Figure 1.4). In fact, even for datasets with little variance, e.g., the

CA dataset (see *Table 2.1*), 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 cycle is crucial. 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.

2.5.3.2 Simulation

The State-of-the-Art approach, i.e., the technique described in [KPS04, KHSR06, KHZS07], to which we refer as *baseline* as we use it as baseline algorithm for comparison, attempts to compute the supported performance level and required battery size based on a representative 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 *Section 2.5.1*, we find the battery capacity and supported performance level with the *baseline* approach from [KHZS07] and one year of calibration data for all three datasets listed in *Table 2.1* and four panel sizes, i.e., 5cm^2 , 10cm^2 , 15cm^2 , and 20cm^2 . The respective performance levels obtained are then used as input to our model to compute the minimum battery capacities required for each dataset. Once these quantities have been found for the *baseline* and our approach (as discussed in *Section 2.4*), we run a simulation with the remaining 10 years of data.

Note that our approach does not necessarily require calibration data. However, since the *baseline* approach relies on a representative power trace, we allow our approach to extract the weather conditions from the calibration data to compute Ω (see *Section 2.3.1*). We use $\tau = [0.1, 0.4, 0.7, 1]$ and $N = 4$ in *Equation (2.2)* and let δ_i be the days with more than 75%, 50%, 25%, and 0% of the maximum expected energy, i.e., the energy obtained with $E_{astro}(\cdot, \Omega = 0.1)$. 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.

2.5.3.3 Results

The results obtained from simulation are summarized in *Table 2.3*, and discussed in the following. As is evident from the table, the *baseline*

approach achieves an acceptable performance only for the CA dataset. As defined in *Section 2.5.1*, acceptable performance means that the battery can support the user-specified performance level indefinitely.

While the *baseline* 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 (see *Table 2.1* for details), the *baseline* approach 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 *baseline* approach are located in Southern California and used locally measured data for design and verification of their approach. From *Table 2.1* 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 *baseline* approach yields a battery size that is about 1.48 times the size of the absolute minimum required battery capacity. For the other two panel sizes, the *baseline* algorithm overestimates the absolute minimum 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 the *baseline*, 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 the *baseline*. This is an important 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 *baseline* model to assume an overly pessimistic negative deviation from ρ_s , and therefore yield a larger battery capacity than necessary.

MI and ON Datasets. The results in *Table 2.3* show that the *baseline* approach does not achieve satisfactory performance for the datasets from Michigan and Ontario (MI and ON, see *Table 2.3*). 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-voltage disconnects, and associated performance penalties.

Tab. 2.3: Expected, and actual power level achieved (in mW), and required battery capacity (i.e., $B = \eta_{out} \cdot B_{nom}$, in Wh) obtained from simulation with 4 panel sizes (in m^2) and the 3 datasets listed in Table 2.1 for the *baseline* approach [KHZS07] and the long-term capacity planning (LT-CP) method proposed in this chapter.

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

When given the opportunity to analyze two full years of calibration data, the *baseline* 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 the *baseline*'s performance depends on the closeness of ρ_s to actual average generation, longer data traces are expected to improve its performance. Interestingly, the *baseline* approach 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, the *baseline* and our approach yield comparable capacities, despite giving the *baseline* the advantage of analyzing two years of calibration data, while our method only used one year of calibration data.

Energy Approximation. In this chapter we focus 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 copes with source variations that lead to an energy deficit. *Figure 2.4* 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 deficit, i.e., on the interval $[d_1, d_2]$, on the right. As is evident, with the largest panel size used ($A_{pv} = 0.02$), the 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 $0.005m^2$ and $0.01m^2$ are about 85% for CA and MI, and around 78% for ON.

For the periods of deficit, i.e., the time periods during which the energy consumption is less than the generation (see *Figure 2.2*), shown on the right hand graphs in *Figure 2.4*, 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 over-provisioned, and will be able to safely bridge short periods that exhibit harvesting opportunities well below the modeled long-term expectations.

2.6 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 represent the actual deployment setup. *Technology parameters* (see *Figure 2.1*), i.e., the




Fig. 2.4: (left) Ratio of total energy $E_{pv}(\cdot)$ (see Equation (2.3)) approximated with our approach, and actual energy input ($E_{real}(d)$, as given by the dataset) for each year, and (right) for each year's period of deficit, i.e., the time during which the generation is below the consumption, see Figure 2.2. Note the scales.




Fig. 2.5: Achieved performance level for CA (left), MI (middle), and ON (right) datasets and four panel sizes with the energy input down-scaled by scaling factor S .

energy conversion and storage efficiency factors, and panel parameters, i.e., size and maximum power output are not discussed, as they linearly scale the observed energy input and energy consumption. Values for these parameters can be obtained from datasheets, or, if not available, from profiling the system.

2.6.1 Environmental Parameter

As is evident from *Section 2.5.2*, 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 effects can 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 *Equation (2.4)*), and simulate the system as discussed in *Section 2.5.1*. The achieved mean duty-cycle over ten years for the three datasets are shown in *Figure 2.5*. The results show that the supported duty-cycle can be maintained without interruption, as long as the *observed* environmental conditions (as defined by Ω , see *Section 2.3.1*) are at least 75%, 85%, and 90% of the *expected* environmental conditions for the CA, ON, and MI datasets respectively.

2.6.2 Setup Parameters

The effects of varying the setup parameters, i.e. inclination and orientation angles, on the harvestable energy are illustrated in *Figure 2.6a*. 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 2.1* 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, *Figure 2.6b* 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 aiming at enabling uninterrupted long-term operation at a predefined performance level.




Fig. 2.6: (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° ($\tau = 0.1$, $k = 0.3$, $R = 0.3$, $\eta_{pv} = 0.1$).

2.7 Chapter Summary

In this chapter we presented a novel approach to systematic 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 astronomical model, which serves as the energy availability model. We further compared our approach to the State-of-the-art [KHZS07] 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 [KHZS07]) 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, *Chapters 3 and 4* extend this work with low-complexity energy harvesting aware dynamic duty-cycling schemes.

3

Dynamic Power Management for Long-Term Operation

3.1 Introduction

In the previous chapter we addressed the problem of power subsystem capacity planning for solar energy harvesting embedded systems. We presented a systematic approach that can, given deployment and technology parameters, aid the designer in finding an appropriate battery capacity for a given application scenario without requiring representative power traces. However, we also noted that, if the conditions observed at the deployment site deviate too significantly from the expectations, a design-time approach alone may not suffice to achieve uninterrupted operation. On the one hand, if the input is below a certain percentage of the modeled expectation, the system may risk premature depletion of the battery because it is unaware of the deviations from the model. On the other hand, if the energy generation by the panel exceeds the expectations, precious energy will go unused that otherwise could have been leveraged by a dynamic approach to improve the system utility. Therefore, to ensure that the system can efficiently use the available energy and operate without interruption over time periods on the order of multiple years, even in the presence of large deviations from the modeled energy input expectations, a mechanism to dynamically adapt the system consumption at runtime is required. Such a mechanism is provided by a so called *dynamic power management* scheme that adjusts the system performance, and therefore its energy consumption, in response to deviations from the expected conditions such that the above objective may be satisfied.

In this chapter we turn our attention to enabling long-term energy neutral operation for solar energy harvesting systems, and apply the insights gained from the development of the capacity planning algorithm to devise a *harvesting-aware dynamic power management* scheme. In the context of long-term application scenarios, e.g., [VKR⁺05, DEM⁺10, CB10, WBB⁺13], this scheme's objective is to compute a system duty-cycle as high as possible that can be sustained. More specifically, the desired *dynamic power management* scheme maximizes the minimum achievable duty-cycle, i.e., performance level, at which the system can operate without interruption over long time periods.

Challenges

A broad range of application scenarios, e.g., [VKR⁺05, DEM⁺10, CB10, WBB⁺13], benefit from a minimum supported performance level that can be sustained over time periods on the order of multiple years. A system enhanced with energy harvesting capabilities can – in theory – operate indefinitely since the energy store can be replenished periodically. Experience has shown, however, that enhancing a battery powered device with energy harvesting capabilities will by itself neither provide a lower bound on the expected sustainable performance level, nor guarantee uninterrupted long-term operation [TJC08]. The reason for this is the dependence on an uncontrollable energy source [KHZS07], i.e., the sun, which exhibits high short-term fluctuations due to meteorological conditions that are hard to model [Buz08] and difficult to predict [HLG06].

Contemporary power management techniques deal with the highly variable energy harvesting opportunities by dynamically adapting the system's performance level at runtime in an attempt to satisfy *Energy Neutral Operation*, i.e., on average the system consumes only as much energy as the panel can generate. Given *Energy Neutral Operation* as the fundamental requirement for energy harvesting systems to achieve long-term operation, numerous methods that attempt to achieve this objective have been proposed, e.g., [KHZS07, VGB07, PBAR09, LW12, LSB⁺12]. These can be classified as (i) predictive, and (ii) reactive approaches. Predictive approaches, e.g., [KHZS07, PBAR09], attempt to satisfy *Energy Neutral Operation* by predicting the harvestable energy during a future time slot, and adapt the performance level accordingly. However, predicting future meteorological conditions is highly complex and may be computationally prohibitive [Buz08]. Therefore, acceptable prediction accuracy with the limited computational resources available on contemporary motes has so far only been possible for short prediction windows, i.e., on the order of minutes to hours. Reactive approaches, on the other hand, attempt to satisfy energy neutrality by scheduling the performance level in response

to changes in the source. This can be done by measuring the energy generation directly, or, as is commonly done, through monitoring the battery fill-level [VGB07], or super-capacitor voltage [LSB⁺12]. The performance of a storage-reactive approach strongly relies on the accuracy of the battery State-of-Charge approximation, and tend to exhibit high duty-cycle variance. Duty-cycle variance is an important consideration, e.g., for surveillance applications, where the system should be available with equal probability at any given point in time [HVV⁺06].

In contrast to previously proposed techniques, we take a radically different approach. In the previous chapter we argued that the source's *long-term* dynamics must be considered for dimensioning the power subsystem, i.e., battery and solar panel, such that short-term fluctuations may be absorbed. Here we will show, that, when aiming at dependable long-term operation, the same holds true for the *dynamic power management* scheme. Rather than attempting to predict or react to the source's *short-term* variations, as the above approaches do, we show that considering the source's *long-term* dynamics enables us to devise an efficient *dynamic power management* scheme that can compute a stable, long-term sustainable performance level at runtime with little computational overhead. Moreover, our approach does not rely on special purpose hardware, as will be shown in *Chapters 5* and *6*.

Contributions

In this chapter we advance the State-of-the-Art in enabling long-term operation of solar powered embedded systems by making the following contributions.

- Firstly, we present a novel end-to-end solution for enabling long-term *Energy Neutral Operation* for solar energy harvesting systems. Our approach encompasses (i) a power subsystem capacity planning approach, which is fully discussed in *Chapter 2*, and (ii) a *dynamic power management* scheme, which is shown to be able to achieve long-term uninterrupted operation with very low duty-cycle variance.
- Secondly, we evaluate and discuss the advantages of *dynamic power management* over only employing appropriate power subsystem capacity planning.
- Thirdly, we present a sensitivity analysis for all relevant parameters involved in our approach.
- Fourthly, we present a discussion of practical issues that must be considered for implementation on resource constrained systems.

- Finally, through simulation with eleven years of data at three different geographical locations, we show that our algorithm outperforms the State-of-the-Art in energy-predictive [KHZS07, PBAR09], and battery-reactive [VGB07] performance scaling approaches in terms of average sustainable performance level by up to 177%, energy efficiency by up to 184%, and duty-cycle stability by up to three orders of magnitude, while incurring zero downtime, i.e., system availability of 100%.

Roadmap

The rest of this chapter is structured as follows. *Section 3.2* briefly reviews the power subsystem capacity planning approach from *Chapter 2*. The proposed *dynamic power management* scheme is discussed in detail in *Section 3.3*. In *Section 3.4* we evaluate the proposed technique through simulation with eleven years of data for three different locations. The benefits of dynamic power management over capacity planning alone is discussed in *Section 3.5*, and a sensitivity analysis of all relevant model parameters is given in *Section 3.6*. In *Section 3.7* we discuss considerations for practical implementations on resource constrained systems. Finally, *Section 3.8* concludes this work with a summary of key findings.

In *Chapter 6* we will further exemplify the benefits of our *dynamic power management* scheme using an *X-Sense* environmental monitoring system [BBF⁺11] deployed over two years as a case study, and demonstrate that significant improvements in system utility can be achieved without risking downtime due to power outages.

3.2 Capacity Planning for Long-Term Energy Neutral Operation

Rather than modeling the energy source's highly variable short-term dynamics and adjust the performance level accordingly, we propose a long-term energy neutral power management scheme for solar energy harvesting systems. Our approach, illustrated in *Figure 3.1*, first invokes a design-time power subsystem capacity planning algorithm to determine the required battery capacity given a set of input parameters that characterize the system and its environment. The intricate trade-offs between battery capacity, and the system and environmental parameters are fully discussed in *Chapter 2*. This capacity planning approach uses an astronomical model to estimate the long-term energy availability based on the




Fig. 3.1: Process flow for long-term solar energy harvesting capacity planning and dynamic power management. Dashed boxes and arrows represent user inputs. The offline capacity planning algorithm computes the achievable duty-cycle and required battery capacity for the given input parameter set, and the dynamic power management algorithm adjusts the system performance level at runtime according to the observed conditions.

annual solar cycle. Then, at runtime, the algorithm proposed in this chapter dynamically computes the performance level, i.e., duty-cycle, based on an adjusted energy availability model such that long-term energy neutrality can be sustained.

3.2.1 System Architecture, Load Model, and System Utility

In this work we assume a harvest-store-use architecture, as described in [SK11], in which the energy to operate the system is always supplied by the battery. There is no bypass path that allows operating the load directly by the panel when surplus energy is available. The system used for the case study presented in *Chapter 6* represents such a harvest-store-use architecture. 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_{sys_i}) and respective active (P_{on_i}) and inactive (P_{off_i}) power dissipation. The total daily energy required to operate at the expected performance level is then defined by *Equation (3.1)*, where $\gamma = 24$ hours.

$$E_{out}(d) = \gamma \cdot \left[P_{cc} + P_{vr} + \sum_i^M \left(DC_{sys_i} \cdot P_{on_i} + (1 - DC_{sys_i}) \cdot P_{off_i} \right) \right] \quad (3.1)$$

For now, we assume a one-to-one relationship between performance level $DC_{sys}(d)$ and utility of the system U , i.e., $U(DC_{sys}(d)) = DC_{sys}(d)$ [CB10]. We revisit this topic in *Chapter 6*, where we refine the definition of system utility in the context of a real system. Note that we are not concerned with how the energy is scheduled and consumed over the course of the day, but rather provide information about disposable energy to an application specific task scheduler. Details on local scheduling of the available energy, and network-wide balancing of the energy budget by changing the communication and/or sensing patterns are beyond the scope of this work, as they are highly application specific. For example, a scheduler's primary focus may be planning the available energy such that a minimum level of operation may be sustained [MBTB07]. Any excess energy may then be used to improve sensing, processing or communication.

3.2.2 Harvesting Conditioned Energy Availability Model

A crucial step in capacity planning consists of estimating the theoretically harvestable energy at a specific point in space and time. *Figure 3.2*




Fig. 3.2: Exemplary solar energy profile. Surplus energy generated by the panel is indicated with the hatched area; the energy deficit is shown by the cross-hatched area.

illustrates the amount of solar energy that can be harvested at a particular geographical location and given harvesting configuration. The figure shows the total daily energy input $E_{real}(d)$ at the end of each calendar day d , and illustrates the high short-term (day-to-day) variability and long-term periodicity (year-to-year) of the source. Also shown is the modeled total expected harvestable energy $E_{in}(d)$ on calendar day d such that true energy conditions are closely approximated, i.e., Equation (3.2) holds where N is the number of days.

$$\sum_{d=1}^N E_{in}(d) \cong \sum_{d=1}^N E_{real}(d), \quad N \gg 1 \quad (3.2)$$

The method to compute $E_{in}(d)$ is based on a simplified astronomical model to estimate the theoretical solar radiation $E_{astro}(t, d, L, \theta_p, \phi_p, \Omega)$. It is parameterized by the time t in hours of calendar day d , the intended deployment site's latitude L , and the panel's orientation and inclination angles ϕ_p and θ_p , respectively. Finally, the environmental parameter Ω represents the expected average meteorological conditions. This is the only unknown input parameter, and can be approximated as described in Section 2.3.1. Although not absolutely necessary, the availability of solar maps or solar energy traces can improve the approximation of the parameter Ω .

Since we are concerned with electrical, as opposed to solar energy, the output of $E_{astro}(\cdot)$ must be conditioned by the technology parameters in Figure 3.1. These specify the panel's surface area A_{pv} , conversion efficiency η_{pv} , battery inherent charging inefficiency η_{in} , and self-consumption and efficiency factors for supervisory and power conditioning circuitry, e.g., battery charge controller efficiency η_{cc} , and consumption P_{cc} . The maximum rated power output of the panel P_{pv} is used to evaluate the maximum energy E_{pv} generated during one hour. Then, with

the above parameters specified, the total electrical energy that can be harvested on calendar day d is approximated with *Equation (3.3)*.

$$E_{in}(d) = A_{pv} \cdot \eta_{cc} \cdot \eta_{pv} \cdot \eta_{in} \cdot \sum_{t=0}^{23hr} \min \{E_{pv}, E_{astro}(t, d, \dots)\} \quad (3.3)$$

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

3.2.3 Power Subsystem Dimensioning

In this section we review the process of computing the power subsystem capacity using the energy availability model such that energy neutral operation over the source's seasonal cycle, i.e., one year, may be achieved. At this point we assume a perfect battery, i.e., no inefficiencies. For a discussion including various battery inefficiencies, the reader is referred to *Chapter 2*.

For the purpose of power subsystem capacity planning we assume a constant daily energy demand $E_{out}(d)$ that must be met. Note that we explicitly keep the dependence on calendar day d , since the energy consumption at runtime varies with the dynamically chosen daily duty-cycle (see *Section 3.3*). Referring to *Figure 3.2*, we observe that the intersections between the energy consumption $E_{out}(d)$ and approximated energy input $E_{in}(d)$ partition the annual solar cycle into time regions of energy surplus, i.e., $E_{in}(d) > E_{out}(d) \forall d \in [d_0, d_1)$, and energy deficit, i.e., $E_{in}(d) < E_{out}(d) \forall d \in [d_1, d_2)$.

According to the model assumptions, the minimum battery capacity B required to support the system during periods of energy deficit is indicated with the cross-hatched area in *Figure 3.2*, and formally stated in *Equation (3.4)*. The first term on the left-hand side defines the amount of energy that is necessary to support the system operation, while the second term represents the expected energy input. The difference is then the minimum required battery capacity.

$$\sum_{d_1}^{d_2} (E_{out}(d) - E_{in}(d)) = B \quad (3.4)$$

In order to achieve uninterrupted operation over multiple years, it is not sufficient to only provision the battery for the period of deficit. The panel must be able to generate enough energy to recharge the battery in addition to what is required in order to sustain operation during periods

of energy surplus, i.e., $d \in [d_0, d_1]$. The constraint on energy generation by the panel is given in *Equation (3.5)*.

$$\sum_{d_0}^{d_1} (E_{in}(d) - E_{out}(d)) \geq B \quad (3.5)$$

The required battery capacity B can then be obtained by varying the performance level (i.e., $DC_{sys}(d)$) and/or the panel area A_{pv} such that *Equations (3.4) and (3.5)* hold.

3.3 Dynamic Power Management

In the previous section we described the design-time energy availability model and power subsystem capacity planning based on the long-term characteristics of the energy source. If we can assume that the design-time model reflects the conditions at the deployment location to within some bounds, the system will be able to run at the performance level for which the power subsystem was designed [BSBT14b]. However, in practice significant deviations from the model must be expected. Such deviations may be caused by transient phenomena, e.g., snow cover and foliage, or persistent occlusions due to trees and buildings. In this section we propose a dynamic power management scheme that can adapt to deviations from the modeled assumptions by dynamically scaling the system performance level, and by doing so enable long-term *Energy Neutral Operation*.

3.3.1 Dynamic Performance Scaling

As discussed in *Section 3.2.3*, in order to achieve long-term energy neutrality, the two constraints from *Equations (3.4) and (3.5)* must be satisfied. The constraint in *Equation (3.4)* states that the battery must be able to supply the difference in energy consumption and generation during periods of energy deficit, i.e., $d \in [d_1, d_2]$ (as shown in *Figure 3.2*). The second constraint states that, in order to ensure that the battery can be fully recharged during periods of energy surplus ($d \in [d_0, d_1]$), the panel must generate energy in excess of what is required to sustain short-term operation. To satisfy these two constraints, we leverage the offline energy model to determine the sustainable system performance level.

To exemplify our approach we consider a concrete example as illustrated in *Figure 3.3*. Without loss of generality, we assume that the design-time model $E_{in}(d)$, which was used to obtain the battery capacity B given panel size A_{pv} , over-estimates the actual energy conditions $E_{real}(d)$.

For simplicity we ignore battery inefficiencies in this discussion, but note that *Algorithm 1* and the evaluation in *Section 3.4* account for these effects. In the following we consider the end of day d and wish to compute the duty-cycle for the entire day $d + 1$ such that long-term energy neutrality may be sustained.

To react to deviations from the modeled energy expectation, we first need to adjust the design-time energy model $E_{in}(d)$ at runtime according to observed conditions. For this purpose, we define the model adjustment factor α in *Equation (3.6)* to scale $E_{in}(d)$, i.e., $\widehat{E}_{in}(d) = \alpha E_{in}(d)$, $\forall d$. The adjustment factor depends on the history window size W in days, which is used to tune the duty-cycle stability. The choice of W has a direct impact on the system's responsiveness to variations in the energy profile, and therefore imposes a trade-off decision between duty-cycle stability and achievable performance level. The effects of the choice of the history window size W are discussed in *Section 3.6.3*.

$$\alpha = \frac{\sum_{d-W}^d E_{real}(d)}{\sum_{d-W}^d E_{in}(d)}, \quad 0 < W \leq d \quad (3.6)$$

Then, referring to *Figure 3.3*, it is evident that, given B and the adjusted energy model $\widehat{E}_{in}(d)$, the modeled consumption $E_{out}(d)$ may not be sustained. For example, a battery capacity dimensioned for bridging the period starting at d_1^* and ending at d_2^* , instead of d_1 and d_2 would be necessary to support $E_{out}(d)$ in *Figure 3.3*. Therefore, to fully, but safely leverage the available battery capacity given $\widehat{E}_{in}(d) < E_{in}(d) \forall d$, we need to find the energy consumption $\widehat{E}_{out}(d) = DC(d) \cdot P_{sys} \cdot \gamma$, where $\gamma = 24 \text{ hours}$ such that the battery and panel constraints in *Equation (3.7)* hold.

$$\sum_{d'_1}^{d'_2} (\widehat{E}_{out}(d) - \widehat{E}_{in}(d)) \leq B \leq \sum_{d'_0}^{d'_1} (\widehat{E}_{in}(d) - \widehat{E}_{out}(d)) \quad (3.7)$$

The limits of summation in *Equation (3.7)* are unknown and depend on $\widehat{E}_{out}(d)$, the quantity we wish to find. However, since the modeled limits are known, or can be computed at runtime, d'_0 , d'_1 , and d'_2 can be found iteratively in discrete time steps, e.g., days, starting with intervals $D_s = [d_0, d_1]$, $D_d = [d_1, d_2]$, which represent the surplus and deficit regions respectively, and adjusting them according to *Algorithm 1* until *Equation (3.8)* evaluates true. D_s^0 and D_d^0 in *Equation (3.8)* denote the first elements in the intervals D_s and D_d respectively.

$$|D_d| \widehat{E}_{in}(D_d^0) - \sum_{d=D_d^0}^{D_d} \widehat{E}_{in}(d) \leq \sum_{d=D_s^0}^{D_s} \widehat{E}_{in}(d) - |D_s| \widehat{E}_{in}(D_s^0) \leq B \quad (3.8)$$




Fig. 3.3: Example where energy availability model $E_{in}(d)$ overestimates the actual energy conditions $E_{real}(d)$. The performance level for $d_{now} + 1$ is computed at the end of day d_{now} , using $\hat{E}_{in}(d)$ with α computed over the past W days.

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

Input: current day d , observed energy input E_{real} ;

Output: duty-cycle for day $d+1$;

$$\alpha = \frac{\sum_{d-w}^d E_{real}(d)}{\sum_{d-w}^d E_{in}(d)};$$

$$\widehat{E}_{in} = \alpha \cdot \eta_{cc} \cdot \eta_{in} \cdot E_{in};$$

$$d'_1 = d_1; d'_2 = d_2; d'_0 = d_0;$$

$$E_{surplus} = \sum_{d'_0}^{d'_1} \widehat{E}_{in}(d) - (\widehat{E}_{in}(d'_1) \cdot (d'_1 - d'_0 + 1));$$

$$E_{deficit} = (\widehat{E}_{in}(d'_1) \cdot (d'_2 - d'_1 + 1)) - \sum_{d'_1}^{d'_2} \widehat{E}_{in}(d);$$

while $((E_{deficit} \leq E_{surplus}) \ \&\& \ (E_{surplus} \leq B_{nom}))$ **do**

if $\alpha < 1$ **then**

 | $d'_0 = d'_0 + 1; d'_1 = d'_1 - 1; d'_2 = d'_2 + 1;$

else

 | $d'_0 = d'_0 - 1; d'_1 = d'_1 + 1; d'_2 = d'_2 - 1;$

end

$$E_{surplus} = \sum_{d'_0}^{d'_1} \widehat{E}_{in}(d) - (\widehat{E}_{in}(d'_1) \cdot (d'_1 - d'_0 + 1));$$

$$E_{deficit} = (\widehat{E}_{in}(d'_1) \cdot (d'_2 - d'_1 + 1)) - \sum_{d'_1}^{d'_2} \widehat{E}_{in}(d);$$

if $(E_{surplus} < E_{deficit})$ **then**

 | $[d'_0, d'_1, d'_2] = D_{prev};$

$$E_{surplus} = \sum_{d'_0}^{d'_1} \widehat{E}_{in}(d) - (\widehat{E}_{in}(d'_1) \cdot (d'_1 - d'_0 + 1));$$

$$E_{deficit} = (\widehat{E}_{in}(d'_1) \cdot (d'_2 - d'_1 + 1)) - \sum_{d'_1}^{d'_2} \widehat{E}_{in}(d);$$

 | **break**;

end

$D_{prev} = [d'_0, d'_1, d'_2];$

end

$$DC(d+1) = \frac{\min(B, E_{surplus}) + \sum_{d'_1}^{d'_2} \widehat{E}_{in}(d)}{P_{sys} \cdot (d'_2 - d'_1 + 1)};$$

The relation in Equation (3.8) is obtained from Equation (3.7) by noting that, under our model assumptions, the energy generation at the start of the deficit period is equal to the consumption on that day (see Figure 3.3). Since we assume a constant energy consumption (i.e., a stable duty-cycle) is desirable, we substitute $\widehat{E}_{in}(D_d^0)$ and $\widehat{E}_{in}(D_d^0)$ respectively, for $\widehat{E}_{out}(d)$ in Equation (3.7), and replace the summations by multiplications.

Note that the maximum battery size that can be supported, given the energy conditions observed, is limited by the energy that can be harvested

during the surplus period, i.e., $B \leq E_{surplus} = \left(\sum^{D_s} \widehat{E}_{in}(d) \right) - \widehat{E}_{in}(d'_1) \cdot |D_s|$. Taking this limitation into consideration, we can use Equation (3.9) to compute the sustainable performance level for day $d + 1$ at the end of day d . Note, $\gamma = 24$ hours, and $E_{surplus}$ is computed as shown in Algorithm 1.

$$DC(d + 1) = \frac{\min(B, E_{surplus}) + \sum^{D_d} \widehat{E}_{in}(d)}{P_{sys} \cdot \gamma \cdot |D_d|} \quad (3.9)$$

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

3.3.2 Worst Case Energy Conditions

The energy neutral dynamic power management approach presented in the previous section relies on the solar energy profile to compute the system's performance level. Under normal circumstances, this is not a problem, as the duty-cycle is adapted according to the long-term dynamics of the source. However, in the case of a prolonged lack of harvesting opportunities, e.g., due to snow cover, the battery must be dimensioned such that this period can be bridged, and the dynamic power management scheme requires a mechanism to deal with the situation. To the best of our knowledge, no other approach considers this scenario.

In order to provision for such conditions, the duration of the expected worst case period, τ days, can be approximated at design-time, and the battery over-provisioned accordingly. Then, at runtime, the power management scheme switches into an emergency mode. For example, we might over-provision the battery with $\sum^{\tau} (DC_e(d) \cdot P_{sys} \cdot 24 \text{ hours})$, and let the duty-cycle be an exponentially decaying function for those days that are below some threshold E_t , i.e., $DC_e(d) = (DC_{min})^{\frac{d+1}{\tau}} \forall \{d | \widehat{E}_{in}(d) < E_t\}$, where DC_{min} is the minimum acceptable duty-cycle.

Using our approach to capacity planning from Chapter 2, and the dynamic power management approach enhanced with the above emergency provisioning, healthy discharge cycles during normal operation can be achieved, as the emergency store is only used in exceptional situations.

3.4 Evaluation

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

3.4.1 Experimental Setup

3.4.1.1 Baseline Algorithms

We compare our approach through simulation against State-of-the-Art implementations of energy-predictive and battery-reactive approaches. Specifically, we implement the predictive duty-cycling scheme from [KHZS07] with two different energy predictors, i.e., EWMA [KHZS07] and WCMA [PBAR09], and one reactive approach, i.e., ENO-MAX [VGB07]. Note that the method from [PBAR09] only provides an energy prediction algorithm but does not discuss dynamic performance scaling, hence we use the algorithm from [KHZS07] to compute the duty-cycle.

We have selected these particular algorithms for the following reasons. The approach from [KHZS07] achieves very good performance with minimal overhead, and is commonly used as a baseline for comparative analysis, e.g., [PBAR09, VGB07]. It is also one of the few techniques that combines prediction and scheduling for solar harvesting systems. The technique in [PBAR09] has been shown to improve the prediction accuracy, but it has not been investigated if the improvement translates into increased system performance. Finally, the method from [VGB07] is a very well-performing representative of the class of storage-reactive approaches.

3.4.1.2 Methodology and Simulation Input Data

To evaluate and compare the performance of the proposed solution, we simulate a solar energy harvesting system with the power management schemes introduced in *Section 3.4.1.1* and the trace data discussed in the following. The simulation environment has been validated as discussed in *Section 6.4*.

For the simulation input data, we resort to the National Solar Radiation Database¹ (NSRD) to obtain a twelve year dataset containing hourly solar radiation measurements at a single observation point in California

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

Tab. 3.1: 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

(CA), Michigan (MI), and Ontario (ON) respectively. The data for the first year of each dataset is used for calibration, and the remaining eleven years are used for simulation input data. The data traces (see *Table 3.1* for details) from the National Solar Radiation Database are given in $Wh \cdot m^{-2}$ of solar energy incident on a flat surface with zero inclination. Hence, to account for smaller panel sizes, inefficiencies of individual components, and losses in energy conversion and storage during simulation, the data is conditioned with an efficiency of a typical midrange solar panel $\eta_{pv} = 10\%$, orientation angle $\phi_p = 180^\circ$, and inclination angle $\theta_p = 0^\circ$. We evaluated different panel sizes, but the results are comparable, hence we only show and discuss results for a panel with $A_{pv} = 5cm^2$. Finally, we consider battery charging and discharging efficiencies with $\eta_{in} = 0.9$ and $\eta_{out} = 0.7$, which are reasonable efficiency factors for Nickel-Metal Hydrid (NiMH) batteries, see e.g., [KPS04, TJC08, PBAR09, VGB07], and empirically validated in [Ste13]. We assume no charge controller inefficiencies, however, i.e., $\eta_{cc} = 1$.

3.4.1.3 Simulation Details

The capacity planning technique from [KHZS07] is used to obtain the battery capacity B and supported power level P_{sys} at full performance (i.e., $DC = 100\%$) using one year of calibration data for each of the three datasets. We do the same with the capacity planning algorithm from *Section 3.2.3*, but do not provision for emergency situations (see *Section 3.3.2*). The results are shown in *Tables 3.2* and *3.3*, and discussed in *Section 3.4.2*.

For each of the baseline implementations we use the authors' recommended parameters, i.e., $K = 3$, $D = 4$, $\alpha = 0.3$ for WCMA [PBAR09], and $\alpha = 0.5$ for EWMA [KHZS07]. For ENO-MAX [VGB07], we use $\alpha = 1/24$, and $\beta = 0.25$. The authors suggest values between 0.25 and 0.75 for β , where lower values are claimed to improve the duty-cycle stability at the cost of performance. We experimented with different values and noted negligible improvements in performance but noticeable increase in duty-cycle variance with increasing values for β . Finally, due to the hourly

Tab. 3.2: Battery capacities and supported power levels obtained with the *baseline* capacity planning approach from [KHZS07] and our approach (LT-CP) proposed in *Chapter 2* for the three datasets MI, ON, and CA (see *Table 3.1*), and a panel size $A_{pv} = 5\text{cm}^2$.

	MI		ON		CA	
	<i>baseline</i>	LT-CP	<i>baseline</i>	LT-CP	<i>baseline</i>	LT-CP
B [Wh]	42.51	93.6	61.48	98.61	56.96	88.96
P_{sys} [mW]	55.73	57	56.75	47.5	73.17	83

Tab. 3.3: Simulation results for the *baseline* approach and our proposed approach (LT-CP) with batteries and expected performance level (i.e., $DC(d) = 100\% \forall d$) from *Table 3.2*.

	MI		ON		CA	
	<i>baseline</i>	LT-CP	<i>baseline</i>	LT-CP	<i>baseline</i>	LT-CP
Offline [%]	28.13	0	33.4	0	24.66	0
P_{mean} [mW]	40.04	57	37.79	47.5	55.11	83
P_{min} [mW]	0	57	0	47.5	0	83
DC Var. [$\frac{1}{mW^2}$]	53.98	0	62.8	0	31.83	0

values given by the National Solar Radiation Database, we use $N_w = 24$ instead of 48 daily update slots for EWMA, WCMA, and ENO-MAX as is done in the respective original implementation. This results in a slight penalty in prediction accuracy, but significantly reduces computational complexity. Recall that our approach performs only one update per day, i.e., $N_w = 1$.

We assume the battery to be fully charged at the start of the simulation, and simulate a low-voltage disconnect hysteresis of 60%, as commonly enforced by modern charge controllers [BAB13]. This means that, if at any time the battery is fully depleted, the load will be reconnected only once the battery has been recharged to 60% of its capacity.

For the history window size W used by our proposed algorithm, we assume $W = 63$ days for all three datasets. The effects of this parameter are further discussed in *Section 3.6*.

3.4.1.4 Performance Metrics

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

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

Mean Power Level. According to *Section 3.2.1*, the system utility is defined by the achievable duty-cycle. However, since the methods evaluated yield different sustainable power levels (see *Table 3.2*), we can not use the duty-cycle alone as a performance metric. Rather, we report the average power level achieved over all simulation time steps (including overriding zeros due to low-voltage disconnects), i.e., $P_{mean} = mean(DC(d)) \cdot P_{sys}$.

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

Duty-Cycle Variance. We report the duty-cycle variance, normalized by the variance of the energy input $E_{real}(d)$ over all simulation time steps, i.e., $Var_{DC} = \frac{var\{DC(d)\}}{var\{E_{real}(d)\}} \forall d$

Energy Efficiency. To compare the energy efficiency of the algorithms, we report the percentage of total energy input that went unused.

3.4.2 Experimental Results

3.4.2.1 Capacity Planning

The State-of-the-Art capacity planning algorithm discussed in [KHZS07] yields the required battery capacity B and sustainable power level P_{sys} , given an energy input trace representative of the conditions at the intended deployment site. Here, we investigate if P_{sys} can indeed be supported over long time periods by simulating the system equipped with a battery of capacity B as obtained with the capacity planning algorithm from [KHZS07], and running at a fixed, full performance power level P_{sys} , i.e., $DC(d) = 100\% \forall d$. Note that any non-zero duty-cycle could have been chosen, as we aim at validating the technique's ability to maintain the expected performance, i.e., in this case full performance. We selected a duty-cycle of 100% since this represents an extreme case. Ultimately, however, the battery should be provisioned with a reasonable lower bound on the duty-cycle required by the application. The dynamic power management scheme then attempts to maximize the duty-cycle at runtime, while taking the battery capacity and panel constraints into consideration.

The results in *Table 3.3* show that the approach from [KHZS07] does not always support the expected performance. The battery obtained with

the long-term energy neutral capacity planning approach from *Chapter 2*, on the other hand, can sustain the expected performance level over the entire eleven years of simulation time. It is worth stating at this point that while in this case our approach also used one year of trace data to approximate Ω , in contrast to [KHZS07], our scheme does not absolutely depend on the availability of such data.

3.4.2.2 Dynamic Power Management

According to the authors, if the calibration data is representative of the actual conditions, the power level P_{sys} obtained with the capacity planning technique from [KHZS07] should be supported at all times. However, the previous experiment showed that this may not always be the case, clearly demonstrating the need for dynamic power management. We thus evaluate and compare the dynamic power management approach proposed in this work against the performance of the power management techniques from [KHZS07] with EWMA and WCMA [PBAR09] predictors, and ENO-MAX [VGB07]. For these algorithms we assume the power subsystem capacity planning from [KHZS07], while our approach uses our capacity planning approach from *Chapter 2* and reviewed in *Section 3.2.3*. We also analyze the scenario in which the baseline algorithms use batteries obtained with our approach.

For simulation we fix the minimum acceptable duty-cycle at $DC_{min} = 50\%$. While this may seem like an unusually high duty-cycle, recall that the power subsystem is designed such that a power level corresponding to $DC = 100\%$ should be possible, and therefore a reasonable lower bound. We then simulate the different approaches with the energy input traces scaled from 50% to 150% to artificially cause model deviations. The results are shown in *Figure 3.4*, and *Table 3.4* lists the performance results averaged over all simulation runs for a static approach (i.e., no duty-cycling is employed), the three baseline algorithms and our proposed scheme, which we call LT-ENO.

From *Figure 3.4*, and *Table 3.4* it is evident that our approach outperforms the baseline algorithms in all respects, except for a few instances where the performance is comparable to that achieved by ENO-MAX. It is particularly noteworthy that the achieved mean power level is bounded closely by the minimum and maximum power levels respectively, illustrating a low duty-cycle variance. For the baseline algorithms, the achieved minimum power level is at most equal to the minimum acceptable power level, i.e., $P_{min} = P_{sys} \cdot DC_{min}$. This means that for the baseline approaches, the minimum achieved power level follows the minimum acceptable duty-cycle defined at design time. With the proposed approach, however, the minimum achievable duty-cycle follows




Fig. 3.4: Mean power level, minimum power level, energy efficiency, and duty-cycle variance with energy input scaled (in increments of 10%) from 50% to 150% of original magnitude for each dataset and $DC_{min} = 50\%$.

Tab. 3.4: Simulation results averaged over all simulation runs shown in *Figure 3.4* for the three datasets, i.e., MI, ON, CA, and the parameters listed in *Table 3.3*. *Note:* Static refers to capacity planning alone, i.e., no dynamic power management is used. LT-ENO refers to our approach proposed in this chapter.

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

the long-term dynamics of the observed energy profile without requiring the designer to specify DC_{min} . Furthermore, considering long-term instead of short-term dynamics has a direct impact on duty-cycle variance. From *Figure 3.4* and *Table 3.4* it is evident that the duty-cycle variance is orders of magnitude lower than that obtained with any of the baseline algorithms. Achieving high duty-cycle stability over long time periods can be a strong requirement in a broad range of application scenarios, e.g., [VKR⁺05, CB10, WBB⁺13].

We further note that, while our approach achieves 100% availability in all simulation runs, the baseline algorithms suffer from depleted batteries for two of the three datasets. In the worst case, this results in system unavailability of up to 620 days (ENO-MAX with ON dataset scaled by 0.5). This behavior is expected since the baseline algorithms are battery and/or source agnostic and, and, more importantly, as has been shown in *Section 3.4.2.1*, the power subsystem is under-dimensioned. In order to perform a fair analysis, and determine if these algorithms could do better, we next evaluate the baseline algorithms with a power subsystem from

Section 3.2.2.

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

In this section, we have shown through simulation that our approach excels in all five performance metrics defined in *Section 3.4.1.4*. The proposed dynamic power management scheme achieves 100% system availability in simulation with eleven years of trace data for different geographical locations. We have shown that the minimum and mean expected performance level can be sustained even when there are deviations from the design-time model assumptions. Since our algorithm leverages the source's long-term dynamics, a very low duty-cycle variance can be maintained while still achieving highly efficient energy usage.

3.5 Benefits of Dynamic Power Management

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

For this purpose, we performed the same simulation discussed in the previous section, but set a static duty-cycle, as obtained from capacity planning. The results, averaged over all simulation runs are shown in the first column of *Table 3.4*. As is evident, the mean achievable duty-cycle without dynamic power management support is close to what is achieved by our dynamic approach. However, if we consider the minimum achievable duty-cycle, the static approach performs significantly worse than our dynamic approach. This is because the static approach can not adjust the




Fig. 3.5: Benefits of dynamic power management. Comparison of mean and minimum achievable duty-cycle with and without dynamic power management.

performance level in response to deviations from the model, and experiences battery depletions when true conditions are below some percentage of the expected conditions.

This is illustrated in *Figure 3.5*, which shows the minimum and mean duty-cycle achieved with the combination of capacity planning (for an expected duty-cycle $DC_{sys} = 70\%$) and the proposed dynamic power management algorithm over a range of scaling factors by which the energy input was scaled. The same is shown for capacity planning alone, i.e., the static approach. Note that in region (a), i.e., for scaling factors 0.5 to 0.9, the static approach was unable to sustain a non-zero minimum duty-cycle, inferring that the system suffered power outages. In this region, the dynamic power management approach achieves a performance level roughly proportional to the expected duty-cycle scaled by the energy input scaling factor. In other words, the dynamic power management approach trades off performance for ensuring continuous operation at an adjusted minimum expected duty-cycle that is dependent only on the energy input. This feature is a clear benefit for systems that require continuous operation.

Region (b) in *Figure 3.5*, i.e., between 0.9 and 1.2 in this case, is a transition region where the static approach achieves slightly higher minimum performance than the dynamic power management scheme. Nevertheless, it is evident that the dynamic power management enabled system is able to improve the performance level in response to increased energy availability. The lower bound of this region depends on the degree of overestimating true conditions, i.e., an effect of capacity planning, while the upper bound of this region is dependent on the reactivity of the dynamic power management algorithm, and therefore related to the history window size W (see *Section 3.6.3*). Finally, region (c) shows the full potential of dynamic power management. That is, the dynamic approach




Fig. 3.6: Achievable maximum, mean, and minimum duty-cycle (top), and duty-cycle variance (bottom) with inclination angle ϕ_p ranging from 0° to 90° for the CA dataset.

continues to adapt to the surplus energy and increases the performance level accordingly.

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

3.6 Sensitivity Analysis

3.6.1 Energy Profile Periodicity

The proposed dynamic power management scheme assumes a certain periodicity and sinusoidal behavior of the energy source. This is a valid assumption, since the tilt in the earth's axis of rotation will cause different incident angles depending on the earth's orbit around the sun, which has a direct impact on the harvestable energy [DHM75]. Despite assuming a stationary solar harvesting setup (i.e., no tracking capabilities), it is nevertheless possible that the expected sinusoidal behavior fails to appear. For example, a natural, or man-made structure may shade the panel over the course of the year such that the typical peaks and troughs (see *Figure 2.6*) are obscured. Occurrences of these environmental effects are considered extenuating circumstances, and therefore not considered in this work. Nevertheless, in the following, we briefly investigate a similar effect due to panel inclination and orientation angles.

3.6.2 Panel Inclination and Orientation

The proposed approach builds upon an energy availability model with deployment specific input parameters. Here we briefly discuss the effects of orientation angle ϕ_p and inclination angle θ_p of the solar panel. First, the effect of ϕ_p has been considered in *Section 3.4.2*, where it was shown that the approach can handle significant deviations from the expected conditions. The inclination angle, however, changes the shape of the annual solar energy profile (see also *Section 2.6.2*). Hence, to evaluate this effect, we simulate the system with the battery provisioned as before, i.e., $\theta_p = 0^\circ$ for the CA dataset, but vary the panel inclination angle θ_p from 0° to 90° for the simulation. *Figure 3.6* shows the maximum, mean, and minimum achievable duty-cycle for the CA dataset. Note that the minimum allowable duty-cycle is fixed at 1%. The results clearly show that the proposed dynamic power management scheme can adapt to unexpected energy profiles, while maintaining very low duty-cycle variance.

3.6.3 History Window Size

As discussed in *Section 3.3.1*, the design-time energy availability model is scaled according to observed conditions over a history window of size W . When applying a large history window size W , the scaling factor α contains information about environmental conditions W days in the past. Large W values reduce the model's reactivity to significant variations in the present energy profile, which has the potential to threaten uninterrupted operation. On the other hand, a short history window enables the system to react to short-term variations of the source, but at the cost of increased duty-cycle variance.

In order to find a suitable trade-off between achievable performance level and duty-cycle stability, we evaluate the adjusted model's approximation accuracy with different values for W . For this we define the performance metric given in *Equation (3.10)*, which considers the model's approximation accuracy through Mean Absolute Percentage Error (MAPE) [AC92], scaled by the variance of α over a time period of N days.

$$\sigma = \frac{\text{var}(\alpha)}{N} \sum_{i=1}^N \left(\left| \frac{E_{\text{real}}(d_i) - \widehat{E}_{\text{in}}(d_i)}{E_{\text{real}}(d_i)} \right| \right) \quad (3.10)$$

The result for $W \in [7, 14, 21, \dots, 140]$ days, and $N = 365$ days is illustrated in *Figure 3.7*, which shows σ for the three different datasets, normalized by the respective maximum value of σ . As expected, the performance metric σ approaches a value, past which there is diminishing improvement in approximation accuracy or stability with increasing W .




Fig. 3.7: Evaluation of history window size W , normalized to maximum value of σ for the respective dataset.

Intuitively, the optimal value for W is likely dependent on the source characteristics, i.e., the variability of the energy profile, and on the length of the deficit period that the battery must be able to bridge. Based on the results in *Figure 3.7*, a history window of $W = 63$ days for all three datasets is considered a suitable parameterization.

3.7 Practical Considerations

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

3.7.1 Measurement Support

The proposed dynamic power management scheme requires that the system can measure or approximate the total daily harvested energy. This can be accomplished by measuring the power output by the panel, or inferring the harvested energy through battery State-of-Charge information e.g., [BAB13, SKJ13]. The former is the preferred choice, but incurs additional overhead in terms of measurement circuitry and continual processing [BKN02]. In *Chapter 5* we present an accurate online battery State-of-Charge approximation algorithm for energy harvesting systems that does not require special purpose hardware. In *Chapter 6*, we show that it is indeed possible to use the technique discussed herein with the State-of-Charge approximation approach from *Chapter 5* when the system is incapable of directly measuring the energy generated by the solar panel.

3.7.2 Global Time Knowledge

Clearly, the proposed technique requires knowledge of global time in order to determine the current calendar day d . Considering that our

approach achieves long-term energy neutrality, and may therefore operate without interruption (see *Sections 3.4.2 and 6.4*), this is not considered a limitation. There exist numerous time synchronization protocols, e.g., [FZTS11, BPBB12], which can be used to synchronize the individual motes' clock with that of a basestation. The basestation, in turn, can keep its local clock up-to-date with the Network Time Protocol (NTP). Details on clock synchronization, however, is beyond the scope of this work.

3.7.3 Battery Inefficiencies

Batteries are non-ideal storage elements, which suffer from a variety of inefficiencies that are dependent on the specific battery chemistry and load behavior [BKN02]. In our model, *charging* and *discharging* inefficiencies are incorporated through η_{in} and η_{out} respectively, as specified by the system designer (see *Figure 3.1*). *Leakage power* is ignored in this discussion. Considering the periodically recurring recharging opportunities, accounting for leakage is not as crucial as it is for purely battery operated devices. *Temperature*, on the other hand, may impact the battery's apparent capacity [BAB13]. Thus, for deployments that are exposed to low temperatures over extended periods of time, it may be necessary to account for the temporarily reduced battery capacity imposed by temperature effects. Finally, *battery aging* is not likely to be a problem, since batteries are generally rated for a few hundred deep discharge cycles, after which a maximum of 80% of the battery's initial capacity can be recovered [BKN02]. With our approach, the battery experiences only one deep discharge cycle per year, and is therefore expected to outlast the lifetime of other system components, e.g., electronics, mechanical parts, *etc.* Note that the solar panel may also experience degradation. However, it has been shown that this tends to be aesthetic in nature, and does not significantly affect the panel's efficiency [CVS⁺07]. A panel's manufacturer generally gives warranties for 20 years, at which point the panel is expected to generate at least 80% of its rated output.

3.7.4 Algorithm Considerations

The proposed algorithm requires a constant amount of non-volatile memory to maintain W values of $E_{real}(d)$, which are necessary to compute α . Furthermore, the system must be able to compute $E_{in}(d) \forall d \in [1, 365]$ at runtime, or alternatively store 365 values representing $E_{in}(d)$ as a look-up table. The computation time is linear with respect to the number of days for which $E_{in}(d)$ is to be determined. Since our approach considers the source's long-term characteristics, we are not concerned with sub-daily energy fluctuations. Hence, the sustainable performance level, i.e., duty-cycle, for the entire day $d + 1$ is computed only once at the end of day

d. Finally, note that the capacity planning algorithm from *Chapter 2*, and reviewed in *Section 3.2.3* is computed offline and relies only on one unknown parameter, i.e., Ω , which can be approximated easily. Similarly, for the dynamic performance scaling algorithm from *Section 3.3.1*, only the history window size W must be determined (see *Section 3.6.3*).

Note that the algorithm presented in *Section 3.3.1* may be optimized. For example, rather than always starting with the modeled limits d_0 , d_1 , and d_2 , we may store the limits obtained on day d and use those as initial conditions on day $d + 1$. However, for the proof-of-concept implementation discussed herein, we are not concerned with the most efficient way to find the intersection of the two functions.

3.8 Chapter Summary

In this work we have demonstrated that appropriate design-time considerations, together with a novel dynamic power management scheme can indeed enable energy neutral operation of solar energy harvesting systems over time periods on the order of multiple years. The proposed dynamic power management scheme leverages an astronomical energy availability model that is also used to dimension the energy harvesting power subsystem.

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

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

4

Optimal Power Management

4.1 Introduction

In the previous chapter we presented a novel dynamic power management technique for energy harvesting systems based on a heuristic. In this chapter we formally study the energy harvesting problem with the objective of maximizing the minimum energy used, and present a dynamic power management scheme with *guaranteed* minimum energy utilization [BKT15].

In *Chapter 3* we showed that lifetime and performance improvements stem from a combination of appropriate power subsystem planning, e.g., [KHZS07, BSBT14b], and energy harvesting-aware dynamic power management, e.g., [VGB07, PBAR09, BSBT14a]. Despite significant research efforts, particularly in the area of runtime power management, the proposed schemes generally aim at maximizing short-term energy efficiency and operate on a best-effort basis. Achieving dependable operation over time periods on the order of multiple years has shown to be difficult despite leveraging a periodically recurring energy source [TJC08]. The inability to provide minimum performance guarantees have so far prevented solar energy harvesting wireless sensor networks to enter important application domains, e.g., safety-critical [AAA⁺07], and long-term monitoring and surveillance [BBF⁺11], where minimum performance guarantees must be given.

In order to advance energy harvesting wireless sensor network technology to the domain of safety-critical systems, we must address the question on how to guarantee a minimal service despite relying on an

energy source that is highly variable both in time and amount of energy availability. Informally, a system that can sustain a pre-defined minimum performance level without interruption is said to provide a minimal acceptable service. If this service-level is not achieved, the system may not be responsive, which – for safety reasons – must be avoided at all cost. As a secondary objective we would like to maximize the utility of the system by leveraging the energy surplus that may likely be available during summer. This “free” energy can be used to e.g., increase the sensing resolution both in time and precision, perform local processing, communicate more frequently, and, in multi-hop networks, participate in forwarding so to relieve the burden from energy starved or otherwise overloaded nodes.

Challenges

Achieving the above objectives is complicated due to the highly variable, and only partially predictable energy profile. An approach that greedily uses the available energy in an attempt to maximize the utility instead of provisioning for times of deficit may cause system outages at a later time. Not only does this incur high penalties due to low-voltage reconnect hysteresis generally enforced by charge controllers [BSBT14b], but it also violates the minimal service objective. On the other hand, conservative short term usage may lead to low minimum service, and the risk of battery overflows. A battery overflow happens when the battery is full and the surplus energy exceeds the maximum system consumption, significantly reducing the energy efficiency because the excess energy can neither be stored nor consumed. Wasting energy is undesirable as it negatively affects the overall system utility. As will be shown, however, in reality it is practically impossible to always fully utilize all available energy without wastage.

As already discussed in *Chapter 3*, contemporary approaches are generally based on feedback-control, and heuristically attempt to trade-off the above objectives without formally defining the underlying optimization problem to be solved. Furthermore, optimality and superiority of a particular approach very much depend on the application scenario under consideration, and the applied evaluation criteria. Most importantly, however, despite over a decade of research in this domain [KS03], to the best of our knowledge, none of the approaches published so far are able to provide minimum performance guarantees for solar energy harvesting systems.

Approach and Summary of Results

In contrast to previous work, we aim to formally study the problem of energy harvesting with the objective of maximizing the minimum energy used across all time intervals. First, we study the simplified clairvoyant setting wherein the harvested energy for any given time interval is exactly known in advance. With this assumption, we derive an optimal algorithm that computes the amount of energy to be used in each time interval such that the minimum energy used, i.e., use function, across all intervals is maximized. To identify an efficient algorithm to compute the optimal use function, we establish a relation between the energy harvesting problem and the shortest Euclidean path problem. This relation then allows us to use well-studied and efficient algorithms for the clairvoyant energy harvesting problem.

Then we study the more general and realistic problem where the harvested energy is not known exactly; instead we only know a conservative estimate of it. With this assumption, we show how a finite horizon control scheme can be used to adaptively update the use function. We prove that under certain assumptions of the estimate, such a finite horizon scheme is guaranteed to provide a certain minimum energy usage that is better than a non-adaptive scheme.

A natural concern with the finite horizon scheme is whether it can be implemented on resource constrained low-power embedded devices used typically in wireless sensor network setups. We therefore study an approximated scheme where certain use functions are planned offline and represented in a look-up table. At runtime, only a few simple operations with the values in the look-up table must be performed to compute the appropriate use function.

The final point to consider is the large variability in the harvested energy. This variability is not modeled in the algorithm, and instead must be adequately represented in the estimate. To this end, we propose a non-uniform scaling function to conservatively scale down the harvested energy estimation.

We experimentally validate the different proposed theoretical ideas with datasets available in a public database and, in *Chapter 6*, with raw data collected from our own long-term deployment. While the theoretical analysis is presented without taking conversion and storage inefficiencies into account, the experiments consider various efficiency factors to evaluate our approach in a realistic setting. With these experiments we show the following key results.

1. The clairvoyant algorithm computes optimal use functions which satisfy certain necessary conditions, as required by the connection to the Euclidean shortest path problem.

2. When relying on the energy estimation from [BSBT14a], the non-clairvoyant, optimal algorithm can lead to failure states, i.e., the battery drains prematurely. However, with our proposed non-uniformly scaled estimator, the algorithm computes sustainable use functions for datasets from diverse locations.
3. The proposed finite horizon scheme improves the minimum energy usage, and greatly increases the total used energy when compared to State-of-the-Art approaches.
4. We show that the efficient look-up table approach results in a marginal loss in performance with respect to the finite horizon scheme. In addition, the memory overhead of the look-up table is found to be reasonably small (about 500 floats) for various data sets.
5. Finally, we show that the proposed approach significantly outperforms all four State-of-the-Art algorithms that are used as baseline for performance comparison. In fact, at best, our approach reaches to within 9.9% of the performance of the optimal clairvoyant algorithm. At worst, the achieved minimum performance is within 29.5% of the clairvoyant approach.

Roadmap

The remainder of this chapter is structured as follows. *Section 4.2* presents a formal system model and defines the optimality criteria for an optimal controller. Then in *Section 4.3* we discuss optimal control of solar energy harvesting systems when the energy input is known in advance. In *Section 4.4* we consider the realistic scenario, where the energy input is not available a priori, and present a finite horizon control scheme that achieves a guaranteed minimum utilization. In *Section 4.6* we provide details on the parameterization of our approach. Then *Section 4.7* presents an extensive discussion of performance results obtained from simulation. Finally, *Section 4.8* presents a brief summary of this chapter.

4.2 Problem Definition

This section introduces a formal system model including the corresponding power management problem. We will also discuss the reasoning behind the model abstractions and show that they are reasonable for a large class of energy harvesting systems. Further discussions on the

usefulness of the model and necessary adaptations if applied to realistic scenarios are described in *Sections 4.6 and 4.7*, and *Chapter 6*.

4.2.1 System Model

The system model used in this chapter is characterized by the following quantities and relations.

Usually, a control algorithm that influences the energy consumption and service of an application will perform its function not continuously but at multiples of a fixed time interval. Therefore, we will use a *discrete time system model*, i.e., time $t \in \mathbb{N}$. The time difference of one is named a *unit time interval*, it could represent a minute, a day, a week or even a month. The time horizon of interest is some known interval $[0, T]$.

The *energy storage element (battery)* has stored energy $b(t)$ at time t . The maximal capacity of the battery is B . For the formal discussion we suppose a loss-free model, i.e., there are no leakage or store-consume inefficiencies. However, for experimental validation we use a realistic model that considers various inefficiencies (see *Section 4.7*). Leakage is ignored, but could be considered by assuming a virtual energy consumption of the application.

The *harvested* and *used* energy during the time interval $[t, t + 1)$ are denoted as $p(t)$ and $u(t)$, respectively. While the harvested energy $p(t)$ is provided by the environment (conditioned by the *technology parameters*, as described in *Section 2.3.2*), the used energy $u(t)$ is under the influence of the harvesting controller. We suppose that the controller can change parameters and modes of the running applications and the underlying hardware device such that the used energy matches any value $u(t)$. As a typical example, the controller may change the duty-cycle of sensing, communication or computation activities, and/or switch certain device components on or off to adjust the consumed energy for a given time interval. Note that we are interested in providing a minimal service-level *at all times*, hence the duty-cycle will not take arbitrarily low values and so introduce undue overhead. Moreover, we assume a properly provisioned power subsystem, e.g., from *Chapter 2*, which can safely provide the maximum power drawn by the system. A more refined model that takes into account an upper-bound on the energy consumption will be introduced in the experimental section, see *Section 4.7*.

Given the current state of the battery $b(t)$, the harvested energy $p(t)$ and the used energy $u(t)$, the *harvesting system* can be described by

$$b(t + 1) = \min \{b(t) + p(t) - u(t), B\}. \quad (4.1)$$

If $b(t) + p(t) - u(t) < 0$, the system is in a *failure state* at time t , due to the battery underflow. We call a use function $u(t)$ *feasible* only if there are no

consequent failure states in $[0, T)$. In practical scenarios, a battery may not be allowed to drop to a zero energy level, which can be considered by assuming a virtually smaller battery than B for the algorithm discussed in the following.

As described above, the controller can influence the energy consumption of the device by changing $u(t)$. Generally, a lower energy also leads to a lower service, e.g., a lower sampling rate or lower data resolution. This fact will be modeled by defining the *utility* $U(t_1, t_2)$ of the use function in time interval $[t_1, t_2)$ as follows

$$U(t_1, t_2) = \sum_{\tau=t_1}^{t_2-1} \mu(u(\tau)) \quad (4.2)$$

for some strictly concave¹ utility function $\mu : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{>0}$. The concavity constraint leads to a diminishing utility if the application consumes more energy, i.e., additional energy is more important in case of a low energy level than in case of an already high energy level. With μ one can also model that beyond a certain consumed energy value there is no additional device utility.

4.2.2 Optimality Criteria of the Controller

Based on the above system model, we can now define reasonable optimality criteria for the controller. Given the stored energy at time $t = 0$, the required stored energy at end of the time-horizon T , and the harvest function $p(t)$ for all $0 \leq t < T$, we are interested in an optimal use function $u^*(t)$ that satisfies the following conditions:

- **C1:** The system never enters a failure state, i.e., we have $b^*(t) + p(t) - u^*(t) \geq 0$, $\forall 0 \leq t < T$, where $b^*(t)$ denotes the stored energy with harvest function $p(t)$ and use function $u^*(t)$.
- **C2:** There is no feasible use function $u(t)$ with a larger minimal energy, i.e., $\min_{0 \leq t < T} \{u(t)\} \leq \min_{0 \leq t < T} \{u^*(t)\}$ for all feasible $u(t)$. This condition is important as we aim to continuously provide a minimal service as high as possible.

There may be multiple use functions that satisfy the defined conditions above. In this case, we may want to choose the optimal use function $u^*(t)$ that maximizes a secondary objective function, such as the following.

¹A strictly concave function μ satisfies $\mu(\alpha x + (1 - \alpha)y) > \alpha\mu(x) + (1 - \alpha)\mu(y)$ for any $0 \leq \alpha \leq 1, x \in \mathbb{R}_{\geq 0}, y \in \mathbb{R}_{\geq 0}$.

- **C3:** Amongst all use functions $u(t)$ that satisfy **C1** and **C2**, we choose the one that maximizes the total *utility* (see Equation (4.2)), i.e., maximizes $\sum_{t=\{0,1,\dots,T-1\}} \mu(u(t))$.

It is by no means obvious if the above criteria can be satisfied simultaneously, or if they are conflicting objectives. At every unit-step, the controller needs to decide whether to use more energy, which would increase the use function for that unit time interval. On the other hand, it removes energy from the battery, which may cause a battery underflow or a lower use function at a future time interval. Furthermore, it is also not known whether there exists an efficient algorithm to determine the optimal use function $u^*(t)$, if it exists. The final challenge is uncertainty and variability in the harvested energy $p(t)$. Due to this, we cannot restrict our focus to the clairvoyant case where $p(t)$ is exactly known in advance.

4.3 Optimal Control

In this section, we will study the problem of computing $u^*(t)$ when $p(t)$ is known in advance. In other words, we consider a clairvoyant algorithm which has perfect knowledge of the harvested energy that can be expected throughout the interval $[0, T)$. Though not practical, this simplified setting allows us to identify the main theoretical ideas that characterize this problem. Subsequently, in Section 4.4 we will consider the realistic problem when only an estimate of the harvested energy is known.

4.3.1 Necessary Conditions for Optimality

The main result of this subsection can be formulated by the following theorem.

Thm. 4.1: *Given a use function $u^*(t)$ such that the system never enters a failure state. If the following relations hold*

$$\begin{aligned} u^*(s-1) < u^*(s) &\Rightarrow b^*(s) = 0 \\ u^*(t-1) > u^*(t) &\Rightarrow b^*(t) = B \end{aligned} \tag{4.3}$$

then $u^(t)$ is (a) optimal with respect to maximizing the minimal used energy among all use functions, and (b) unique.*

In other words, if we are able to construct a use function that satisfies Equation (4.3) and never leads to a failure, then the minimal used energy is optimally maximized. We refer to this use function $u^*(t)$ as the optimal use function. This result defines a necessary and sufficient condition on the optimality to satisfy criteria **C1** and **C2**. Further, as only one use function satisfies this necessary condition, it also trivially satisfies the criterion **C3**.

We will now show the sequence of lemmas that lead to the above fundamental result.

Lem. 4.1: (Necessary Condition) *Any optimal use function $u^*(t)$ satisfies*

$$\forall s \leq \tau \leq t : 0 < b^*(\tau) < B \Rightarrow \forall s-1 \leq \tau \leq t : u^*(\tau) = u^*(t). \quad (4.4)$$

and

$$\begin{aligned} u^*(s-1) < u^*(s) &\Rightarrow b^*(s) = 0 \\ u^*(t-1) > u^*(t) &\Rightarrow b^*(t) = B. \end{aligned} \quad (4.5)$$

Proof. We first show that

$$\forall s \leq \tau \leq t : 0 < b^*(\tau) < B \Rightarrow \forall s-1 \leq \tau \leq t : u^*(\tau) = u^*(t).$$

We show that a use function $u(\tau)$ that does not satisfy the above criterion can be replaced by a use function $u'(\tau)$ with a higher minimum used energy. The change will be local, i.e., it will only influence two values of $u(\tau)$ in $s-1 \leq \tau \leq t$, i.e., $u(\tau_1)$ and $u(\tau_2)$. Moreover, we will guarantee that $\sum_{\tau=s-1}^t u(\tau) = \sum_{\tau=s-1}^t u'(\tau)$ and therefore, the stored energy functions will satisfy $b'(\tau) = b(\tau)$ for all $0 \leq \tau \leq s-1$ and $t+1 \leq \tau \leq T$. This is achieved by setting $u(\tau_1) + u(\tau_2) = u'(\tau_1) + u'(\tau_2)$.

The change of some values of $u(\tau)$ must not violate the conditions on the stored energy. The maximum allowed variation on $u(\tau_1)$ and $u(\tau_2)$ is $\epsilon = \min_{s \leq \tau \leq t} \{|b(\tau)|, |B - b(\tau)|\} > 0$

As $u(\tau)$ is not constant in $s \leq \tau \leq t$, there exists a maximum and a minimum $u(\tau)$ in $s \leq \tau \leq t$, say $u(\tau_1)$ and $u(\tau_2)$, respectively with $\tau_1 \neq \tau_2$. We now choose

$$u'(\tau_1) = u(\tau_1) - \delta \quad u'(\tau_2) = u(\tau_2) + \delta$$

with

$$\delta = \min \left\{ \frac{u(\tau_1) - u(\tau_2)}{2}, \epsilon \right\} > 0.$$

We finally have to show that this change leads to a higher minimum used energy, and hence $u(\tau)$ was not optimal as assumed.

Now we show that

$$u^*(s-1) < u^*(s) \Rightarrow b^*(s) = 0,$$

$$u^*(t-1) > u^*(t) \Rightarrow b^*(t) = B.$$

From the earlier condition, we know that $u^*(\tau-1) \neq u^*(\tau) \Rightarrow b^*(\tau) = 0 \vee b^*(\tau) = B$. Suppose now that there is some use function with $u(s-1) < u(s)$ and $b(s) = B$. Then we can use the same method as in the proof of the first part of Lemma 4.1 to show that in this case, u was not optimal. To this end, we replace the value of u at $s-1$ by the larger value $u(s-1) + \delta$ and that at s by the smaller value $u(s) - \delta$ by choosing a suitably small positive value of $\delta \leq (u(s) - u(s-1))/2$ which does not lead to a failure or waste state. The stored energy at time s is now $B - \delta$ which is still feasible, but the overall utilization would be larger. Therefore, u was not optimal. The proof for the second relation is analogous. \square

According to this Lemma, the used energy should be constant as long as the battery is neither full nor empty. In addition, if an optimal use function grows (shrinks), then the corresponding battery level must be empty (full). Interestingly, this necessary condition applies for all possible harvested energy functions $p(t)$.

Lem. 4.2: (Uniqueness) *If there exists a use function $u^*(t)$ that satisfies the necessary optimality conditions of Lemma 4.1 and that does not lead to a failure, then it is unique.*

Proof. Let us suppose that there are two different feasible solutions u_1 and u_2 that satisfy the conditions of Lemma 4.1. At time $\tau = 0$, the stored energy is given as $b(0)$, both use functions result in the same given stored energy at the end of interval $b(T)$, and no use function leads to a failure or waste state.

As u_1 and u_2 are different, there is some first time instance $\tau = s$ where both functions are different, e.g. $u_1(s) > u_2(s)$. As the total harvested energy equals the used energy (no waste state) we have

$$\sum_{\tau=0}^{T-1} u_1(\tau) = \sum_{\tau=0}^{T-1} u_2(\tau).$$

Therefore, there must exist a time $\tau = t > s$ where $u_1(t) < u_2(t)$. We will show that such a time instance can not exist under the conditions of the Lemma.

Let us suppose that t is the first time instance with $u_1(t) < u_2(t)$, i.e., $u_1(\tau) \geq u_2(\tau)$ for all $0 \leq \tau \leq t-1$. The stored energy under the




Fig. 4.1: (top) Example of an optimal use function $u^*(t)$ for a given harvest function $p(t)$ and (bottom) the corresponding stored energy $b^*(t)$.

use function u_i is denoted as b_i . From $b_i(\tau + 1) = b_i(\tau) + p(\tau) - u_i(\tau)$ and the existence of time s we find that $b_1(t) < b_2(t)$.

In order to realize such a situation, we either have $u_1(t - 1) > u_1(t)$ or $u_2(t - 1) < u_2(t)$ or both conditions hold. We will show that none of the conditions is feasible. Suppose we have $u_1(t - 1) > u_1(t)$. Because of Lemma 4.1 we find $b_1(t) = B$. But this is not possible as $b_1(t) < b_2(t)$ and $b_2(t) < B$. Suppose we have $u_2(t - 1) < u_2(t)$, then we find $b_2(t) = 0$, again due to Lemma 4.1. But this is not possible as $b_1(t) < b_2(t)$ and $b_1(t) > 0$. \square

The above result is quite surprising but an example may help to interpret it. To this end, we use a data trace of harvested solar energy with a time granularity of one week. Figure 4.1 shows 100 weeks of an optimal use function $u^*(t)$ and the corresponding battery state $b^*(t)$. As is evident, when the battery is neither full or empty, $u^*(t)$ is constant. Further, if $u^*(t)$ decreases, then the battery is full and if $u^*(t)$ increases, then the battery is empty. This seems highly counterintuitive as one would expect that in case of a full battery it is reasonable to increase the energy usage, and in case of an empty battery one would decrease the energy usage. However, considering that the algorithm seeks to fully leverage the available energy (including the energy that is stored in the battery), an empty battery marks the end of a period of deficit, i.e., the period during which the battery is being fully discharged (see Figure 4.1). Therefore, it makes sense to increase the energy usage upon entering such a period as now there is more harvested energy available. Conversely, a full battery marks the

end of a surplus period and therefore, the usage of energy needs to be reduced.

After having shown the conditions for an optimal use function $u^*(t)$ it remains to be shown whether there is an efficient algorithm to determine the same for any given energy input $p(t)$.

4.3.2 Algorithm

In this section we will show that there is an unexpected relation between the optimal power management problem for harvesting systems as defined above and a problem in computational geometry, i.e., the shortest Euclidean path in simple polygons. The geometric problem has been widely studied and efficient algorithms are available. In particular, in [GHL⁺87] an $O(n)$ algorithm is described where n denotes the number of polygon vertices. Hence, it can be expected that the optimal control problem can be solved efficiently as well.

Similar to [CSSJ11], we will first establish a connection between the two seemingly unrelated problems. To this end, we define a (*harvesting*) *polygon* by all points (t, σ) with integer time $t \in [0, T]$ and $\sigma_l(t) \leq \sigma \leq \sigma_u(t)$ where

$$\sigma_l(t) = \sum_{\tau=1}^t p(\tau - 1) \quad \forall 1 \leq t \leq T \quad (4.6)$$

$$\sigma_u(t) = \sigma_l(t) + B \quad (4.7)$$

and $\sigma_l(0) = 0$. In other words, $\sigma_l(t)$ denotes the harvested energy in the time interval $[0, t]$. Therefore, the polygon contains all points (t, σ) where σ is between the energy harvested in $[0, t]$ and this value shifted up by the capacity of the battery B .

Now, a *feasible path* $f(t)$ in the (*harvesting*) polygon starts at $f(0) = B - b(0)$, ends at $f(T) = \sigma_u(T) - b(T)$, is monotonically increasing, and never leaves the polygon, i.e., $\sigma_l(t) \leq f(t) \leq \sigma_u(t)$ for all $t \in [0, T]$. Furthermore, as we consider a discrete time model, f is assumed to be piecewise-linear with discontinuities only at integral time units.

As has been shown in [CSSJ11], there is an one-to-one correspondence between any feasible path $f(t)$ in the harvesting polygon and a feasible use function $u(t)$ with no waste where

$$u(t) = f(t + 1) - f(t) \quad \forall 0 \leq t < T. \quad (4.8)$$

Moreover, the shortest feasible path $f^*(t)$ in the harvesting polygon yields the optimal use function $u^*(t)$ that maximizes the throughput in the network [CSSJ11].

Algorithm 2: Iterative algorithm for optimal power management.

Data: $b(0); b(T); B; \epsilon; T; (\forall 0 \leq t < T : p(t));$
Ensure $(\forall 0 \leq t < T : u(t));$
for $0 < t \leq T$ **do**
 $\sigma_l(t) \leftarrow \sum_{\tau=1}^t p(\tau - 1);$
 $\sigma_u(t) \leftarrow \sigma_l(t) + B;$
end
 $f(0) \leftarrow B - b(0);$
 $f(T) \leftarrow \sigma_u(T) - b(T);$
forall $0 < t < T$ **do**
 $f(t) \leftarrow (\sigma_u(t) + \sigma_l(t))/2;$
end
repeat
 $f' \leftarrow f;$
 for $t = 1$ **to** $T - 1$ **do**
 $f(t) \leftarrow (f(t - 1) + f(t + 1))/2;$
 $f(t) \leftarrow \max[\min\{f(t), \sigma_u(t)\}, \sigma_l(t)];$
 end
until $\max\{|f' - f|\} < \epsilon;$
forall $0 \leq t < T$ **do**
 $u(t) \leftarrow f(t + 1) - f(t);$
end

It appears that the runtime optimal algorithm [GHL⁺87] is difficult to implement due to the use of advanced data structures and complex operations. In *Section 4.5* we will propose a strategy to determine explicit solutions to the optimal power management problem that can be easily stored and used on resource-constrained platforms. Here, we will describe a very efficient iterative algorithm that can be used on resource-rich platforms or server-controlled devices. This approximation algorithm has guaranteed convergence properties in terms of precision vs. runtime and is based on the ideas presented in [LK07], [LK11]. *Algorithm 2* uses the parameter ϵ as a stopping criterion, i.e., the maximal change in any path coordinate from one iteration to the next.

The first lines in *Algorithm 2* define the harvesting polygon as defined in *Equations (4.6)* and *(4.7)* as well as the initial and final point of a feasible path. Then, the initial path is chosen to be in the middle between the upper and lower boundaries of the harvesting polygon. The main iteration attempts to make an as straight as possible path within the polygon, i.e., it puts a point in the middle between its neighbors unless it touches the upper or lower boundary. This algorithm appears to be very efficient; runtime results are described in the experimental *Section 4.7*.

4.4 Finite Horizon Control

The previous section provided conditions and an algorithm for optimal clairvoyant control: If the initial state $b(0)$, the final state $b(T)$ of the battery, as well as the harvested energy $p(t)$ for all time intervals in $[0, T)$ are given, then the optimal and unique use function $u^*(t)$ can be determined. But in reality, only an estimate of the harvested energy in the future can be provided. In this section, we present a finite horizon control scheme that allows to formulate bounds on the achieved optimality despite of the fact that only an approximation of the future energy intake is available.

To this end, we suppose that for all time instances $t \geq 0$ there is an estimated harvested energy function $\tilde{p}(t)$ available that provides a lower bound on the actually harvested energy $p(t)$ in $[t, t + 1)$:

$$p(t) \geq \tilde{p}(t) \quad \forall t \geq 0$$

Due to the fact that we deal with solar energy harvesting systems, we can suppose that $\tilde{p}(t)$ shows periodic behavior, e.g., a yearly repeating summer/winter pattern. We further suppose that the periodicity is P , and therefore we have

$$\tilde{p}(t) = \tilde{p}(t + P) \quad \forall t \geq 0$$

Moreover, the time interval we consider is expected to be an integral multiple of the period, i.e., $T = P, 2P, \dots$

The derivation of the finite horizon control scheme and its correctness proof involves three steps: At first, we will derive the optimal control for the above periodic estimated harvested energy. In the second step, we will use this result to derive the adaptive finite horizon scheme. Finally, we will show that the adaptive finite horizon scheme will never provide worse results than the optimal control for the estimated harvested energy $\tilde{p}(t)$. This way, the three optimality criteria from *Section 4.2.2* can be satisfied, i.e., (a) the system never enters a failure state, (b) a lower bound of the use function can be provided as well as (c) a lower bound on the *utility* of the use function. Note that the specific choice of utility function is very application specific (see also *Section 4.7*).

4.4.1 Optimal Periodic Control

As described above, we suppose a periodic estimated harvested energy function $\tilde{p}(t)$. By assuming that the state of the battery $\tilde{b}(P)$ at the end of the period equals that at the beginning $\tilde{b}(0)$, we find that all essential quantities are periodic in case of the optimal control, i.e., optimal use function, energy harvesting function and battery state: $\tilde{u}^*(t) = \tilde{u}^*(t + P)$ and $\tilde{b}^*(t) = \tilde{b}^*(t + P)$.

Algorithm 3: Iterative algorithm for periodic optimal control.

Data: $P; B; \epsilon; (\forall 0 \leq t < P : \tilde{p}(t))$;
Ensure: $(\forall 0 \leq t < P : \tilde{u}(t)) ; (\forall 0 \leq t < P : \tilde{b}(t))$;
forall $0 < t \leq P$ **do**
 $\sigma_l(t) \leftarrow \sum_{\tau=1}^t \tilde{p}(\tau - 1)$;
 $\sigma_u(t) \leftarrow \sigma_l(t) + B$;
end
 $f(0) = B/2$;
 $f(P) = \sigma_l(P) + B/2$;
forall $0 < t < P$ **do**
 $f(t) \leftarrow (\sigma_u(t) + \sigma_l(t))/2$;
end
repeat
 $f' \leftarrow f$;
 $f(0) \leftarrow (f(P - 1) - \sigma_l(P) + f(1))/2$;
 $f(0) \leftarrow \max[\min\{f(0), B\}, 0]$;
 for $t = 1$ **to** $P - 1$ **do**
 $f(t) \leftarrow (f(t - 1) + f(t + 1))/2$;
 $f(t) \leftarrow \max[\min\{f(t), \sigma_u(t)\}, \sigma_l(t)]$;
 end
 $f(P) \leftarrow f(0) + \sigma_l(P)$;
until $\max\{|f' - f|\} < \epsilon$;
forall $0 \leq t < P$ **do**
 $\tilde{u}(t) \leftarrow f(t + 1) - f(t)$;
 $\tilde{b}(t) \leftarrow \sigma_u(t) - f(t)$;
end

Let us suppose now that the control horizon T equals the period of the harvested energy: $T = P$. The optimized use function can now be determined using a periodic variant of *Algorithm 2* where we assume $\tilde{b}(0) = \tilde{b}(P)$ and we iterate the path update in a circular way with $f(t + P) = f(t) + \sigma_l(P)$, see *Algorithm 3*.

In summary, *Algorithm 3* computes an optimal periodic battery state \tilde{b} and use function \tilde{u} given a periodic energy harvesting function \tilde{p} , its period P and the battery capacity B . Based on this result, we will define the finite horizon scheme in the next subsection.

4.4.2 Finite Horizon Scheme

Following the ideas of receding horizon control, see e.g., [KH06], we replace the unknown harvested energy function $p(t)$ with its estimate $\tilde{p}(t)$ and optimize the future use function assuming the current battery state.

Algorithm 4: Finite horizon control for the optimized control of solar harvesting systems.

Data: $t; b; T; \tilde{b}(t + T); (\forall 0 \leq \tau < T : \tilde{p}(t + \tau)) ;$
Ensure: $u ;$
 $b(0) \leftarrow b ; b(T) \leftarrow \tilde{b}(t + T) ;$
forall $0 \leq \tau < T$ **do**
 $p(\tau) \leftarrow \tilde{p}(t + \tau) ;$
end
execute *Algorithm 2* with the above inputs and the result
 $(\forall 0 \leq t < T : u(t)) ;$
 $u \leftarrow u(0) ;$

From this optimized use function, we only execute the first time step. After this time step, we again optimize the future use function, and so on.

Algorithm 4 implements this strategy, i.e., at each time step t it uses the current battery state $b(t)$ and provides the optimized use function $u(t)$. The algorithm requires the following static input data: the estimated energy function $\tilde{p}(t)$, the optimized battery value $\tilde{b}(t + T)$ from *Algorithm 3* and the length of the finite horizon T . Note that \tilde{p} and \tilde{b} are periodic in P , i.e., $\tilde{p}(t) = \tilde{p}(t + P)$ and $\tilde{b}(t) = \tilde{b}(t + P)$.

The performance of the finite horizon scheme crucially depends on the accuracy of the estimate $\tilde{p}(t)$. If we over-estimate the actual harvested energy, then it is possible that *Algorithm 4* computes an aggressive use function which leads to a failure state, i.e., the battery drains out. To avoid this we must have a conservative estimate where $p(t) \geq \tilde{p}(t)$. In the next subsection we will identify what guarantees can be provided with such an estimator.

4.4.3 Guarantees

The results of this subsection are presented assuming that the estimate $\tilde{p}(t)$ is *static*. By this we mean that the estimate for the amount of harvested energy in any given time interval within the period is a constant across all periods. For instance, for the solar data, we assume that the estimate of the harvested energy in a certain week is the same across all years. As we illustrate in *Section 4.7*, identifying and using such a constant estimate based on an astronomical model is effective for datasets from several locations. Further, we consider a conservative estimate, i.e., the estimate at each time step does not exceed the actual harvested energy: $\tilde{p}(t) \leq p(t)$.

With a static estimate $\tilde{p}(t)$, there are two options for designing a non-clairvoyant algorithm. First, with *Algorithm 2* we can use the estimate $\tilde{p}(t)$ to compute a static use function $\tilde{u}(t)$ for time steps within the periodic

window $[0, P)$. Second, with the finite horizon scheme from *Algorithm 4* we can recompute the use function $\tilde{u}(t)$ at each time step. The general expectation is that the finite horizon scheme will perform better than the static scheme. Can we formally establish such a relation? In other words, can we prove that the minimum use function under the finite horizon scheme is not less than that of the static scheme? Such a result will enable us to guarantee performance with respect to the estimate $\tilde{p}(t)$.

In *Algorithm 4*, we execute a clairvoyant optimal control at every time step t . Let us suppose that the initial battery at time $t = 0$ equals $\tilde{b}(0)$ as determined by *Algorithm 3*. Then the first control step is $u(0) = \tilde{u}(0)$. If the actual harvested energy is larger than or equal to the estimated one, the initial battery state for the next time step is larger than the estimated $\tilde{u}(1)$. Suppose that the optimal use function satisfies some monotonicity, then we will never have a smaller use function value than in the case of the estimated harvested energy. Such a monotonicity should state that the optimal use function can never increase with (a) increasing initial battery state and (b) increasing energy harvesting function.

We now show this monotonicity property and based on this, the guarantees we can provide for the finite horizon control algorithm.

Thm. 4.2: *Given an initial and final battery state $b(0)$ and $b(T)$, respectively, the battery capacity B , and the harvested energy function $p(t)$ for all $0 \leq t < T$. Then all values $u^*(t)$ for $0 \leq t < T$ of the optimal use function $u^*(t)$ according to Theorem 4.1 are monotonically increasing with increasing $b(0)$, B , $p(t)$ for $0 \leq t < T$, and with decreasing $b(T)$.*

Proof. We use a constructive proof principle. We are given the optimal use function $u^*(t)$ before applying any change to $(b(0), B, p(t)$ for $0 \leq t < T, b(T))$. Then we can construct an adapted feasible use function for the new input parameters that has no smaller use function values than $u^*(t)$. The optimal use function for the changed input parameters can only be better.

If $b(0)$ increases by some value Δ , then we replace $u^*(0)$ by $u^*(0) + \Delta$. All other values $u^*(t)$, $t > 0$, remain unchanged. If $b(T)$ decreases by some value Δ , then we replace $u^*(T - 1)$ by $u^*(T - 1) + \Delta$. If B increases by some value, then $u^*(t)$ is still feasible, though not optimal in general. If $p(t)$ increases for some t by some value Δ , then we replace $u^*(t)$ by $u^*(t) + \Delta$ and the state of battery at time $t + 1$ remains the same as before. \square

Based on this property, we can now proof the main result for the finite horizon control algorithm.

Thm. 4.3: *Given a periodic estimated energy harvesting function $\tilde{p}(t)$ with period P , a battery with capacity B , the corresponding optimal periodic use function $\tilde{u}^*(t)$ and the corresponding periodic battery state $\tilde{b}^*(t)$. Furthermore, given an energy harvesting system with the same capacity B but with an initial battery $b(0) \geq \tilde{b}(0)$ and an actual energy harvesting function $p(t) \geq \tilde{p}(t)$ for all $t \geq 0$. Then the following holds: If we execute the finite horizon control according to Algorithm 4 for each time step t , then the resulting use function satisfies $u(t) \geq \tilde{u}^*(t)$ for all $t \geq 0$.*

Proof. At every time step t when we execute Algorithm 4, we can compare the situation to the optimal periodic case with the estimated energy harvesting function. We use an inductive argument. Initially, the battery state is larger than that of the optimal periodic case as $b(0) \geq \tilde{b}(0)$. Suppose now we are at time t and the current battery state satisfies $b(t) \geq \tilde{b}(t)$. Then we will show that the determined use function is larger or equal than $\tilde{u}^*(t)$ and that $b(t+1) \geq \tilde{b}(t+1)$. The finite horizon control determines at time t a use function based on the horizon $[t, t+T]$ and the estimated harvested energy in this interval. As the final battery state at time $t+T$ equals that of the optimal periodic case and as the initial battery state is higher, the determined use function is larger than $\tilde{u}^*(t)$. The battery state at time $t+1$ is larger or equal due to two independent reasons. At first, the actual harvested energy in $[t, t+1)$ is larger or equal than the estimated one. Moreover, a larger initial battery at t leads to a larger battery at the next time step $t+1$ even in the case of equal harvested energy due to the triangle inequality in the equivalent shortest path problem. \square

In other words, the use function when executing the finite horizon control algorithm is always larger than the optimal periodic one, e.g., as computed using Algorithm 3. Therefore, the minimal value of the use function as well as its utility are lower bounds for the online horizon control. This way, under the assumptions mentioned in Theorem 4.3 we can provide the following off-line guarantees: The system will never enter a failure state, and the minimal use function value as well as the utility values are larger than those from Algorithm 3.

4.5 Implementation

A natural concern with the proposed finite horizon control is the need to execute the clairvoyant algorithm for each time-step. This can be challenging on resource-constrained embedded processors of typical WSN motes. As a response to this challenge, in this section we study the possibility to

Tab. 4.1: Non-volatile memory requirements for the LUT approach (in floats). See Table 4.2 for details about the eight datasets.

Dataset								Mean
TX2	TX1	CA	MD	MI	OR	ON	AK	515
454	494	478	555	542	580	507	509	

offload the computing complexity to an offline phase and export only a small look-up table which is then used online. In Section 4.7, we quantitatively study the loss in optimality due to this approximate solution for real data-sets.

Under the assumption of a static estimate of the harvested energy, we can consider the finite horizon scheme as identifying the current use function $u(t)$ as an explicit function F of the current time-step t (with respect to the period) and the current battery state $b(t)$. In this sense, we can offline pre-compute this explicit function F . At runtime, we sense the battery state and read off the use function from F .

For a given time-step, $F(t, \cdot)$ is a function of only the battery. We consider approximating $F(t, \cdot)$ by a piecewise-linear function \bar{F}_t . The function \bar{F}_t can be computed offline for each possible battery state, discretized by a certain step. In our experiments, we discretize the battery by steps of 1%. Note that the discretization has a direct effect on the LUT size. However, considering that sub-percentage battery State-of-Charge approximation accuracies are unlikely (see Chapter 5), this is considered a reasonable lower bound on discretization steps. Therefore, the values given in Table 4.1 are an upper bound on the storage requirements. For each battery state, we explicitly use Algorithm 2 to compute the optimal clairvoyant use function and thereby identify the use function for the first time-step. These linear functions can then be represented concisely in a look-up table. The non-volatile storage requirements of these tables for the datasets used are listed in Table 4.1. As a result of using a look-up table, the corresponding online algorithm is reduced to very few operations, as shown in Equation (4.9), where b_d refers to the current battery fill level, and $c(\cdot)$ and $f(\cdot)$ refer to ceiling and floor functions respectively, which are assumed to return integer values.

$$u(d) = LUT[d, f(b_d)] \cdot (b_d - f(b_d)) + LUT[d, c(b_d)] \cdot (c(b_d) - b_d) \quad (4.9)$$

The major concern of this approach is the loss in optimality due to the approximation. In Section 4.7 we quantify the marginal loss, and

demonstrate that such a look-up table is an effective and efficient solution with very little runtime overhead.

4.6 Algorithm Parameterization

The control scheme introduced in the previous section requires an estimation of the harvestable energy for each unit time interval to compute the optimal use function. While many different estimation approaches may be possible, we opted to use an astronomical model [DHM75] as the basis for the energy estimation model. In this section we briefly review the astronomical model that serves as a dynamic energy estimator, and discuss parameter selection for the proposed control scheme.

4.6.1 Energy Estimation Model

To estimate the total solar energy available at a particular geographical location without requiring trace data, we leverage the energy estimator discussed in detail in *Section 3.2.2*. It has been shown to very accurately approximate the average harvestable energy for a given geographical location and time of year [BSBT14b, TJC08].

4.6.2 Horizon of Control

With the objective of achieving long-term minimum performance guarantees, it is reasonable to consider the annual, rather than diurnal solar cycle for the control horizon. This means that at any time instance τ , the algorithm from *Section 4.4.2* is given $\tilde{p}(t) \forall t \in [\tau, \tau + T]$ as estimation input, where $T = 1$ year.

The algorithm further requires current and expected final battery fill levels as input, i.e., $b(\tau)$ and $b(\tau + T)$, to compute the optimal energy allocation. We assume the battery to be initially charged to 50% of its capacity, i.e., $b(0) = B/2$, and the expected final fill-level to be equal to the optimal fill-level, i.e., $b(\tau + T) = \tilde{b}(\tau)$. The optimal fill-level $\tilde{b}(t)$ is found by executing Algorithm 3 with the energy model from the previous section as input.




Fig. 4.2: (top) Average achieved minimum utilization and (bottom) utility for all datasets (normalized to maximum values). Note the different scales. Errorbars represent the maximum deviation from the mean for all datasets.

4.6.3 Frequency of Control

With the energy estimation model and the control horizon established, we now investigate the necessary and appropriate duration of the *unit time interval* t for which the algorithm is to compute the supported energy utilization $u(t)$. It is generally desirable to keep the update intervals long so to reduce the algorithm's computational overhead and keep a stable service-level. However, in the case of solar energy harvesting, meteorological conditions, which are known to be difficult to predict [HLG06], are the major contributor to variations from the astronomical model expectations. Hence, the high variability in the energy profile may necessitate shorter intervals, so to allow the control algorithm to react to model deviations.

In order to find the appropriate update interval, we simulate our algorithm (see *Section 4.7* for simulation details) with $t \in [1, 2, 4, 6, 8]$ weeks and all datasets listed in *Table 4.2*. The resulting average minimum utilization and utility achieved are shown in *Figure 4.2* (both normalized to their respective maximum values). As is to be expected, shorter update periods allow the algorithm to make most efficient use of the available energy, while increasing t adversely impacts both the utility and minimum due to somewhat conservative estimation.

With the goal of enabling competitive long-term minimum performance guarantees, we opt for a unit time interval of $t = 1$ week, which is considerably longer than the baseline algorithms that we use to benchmark our approach in the next section.

4.7 Evaluation

4.7.1 Experimental Setup

4.7.1.1 Baseline Algorithms

We compare our approach against a number of State-of-the-Art implementations of energy-predictive and battery-reactive power management approaches. Specifically, we implement the predictive duty-cycling scheme from [KHZS07] with two different energy predictors, i.e., EWMA [KHZS07] and WCMA [PBAR09], and one reactive approach, i.e., ENO-MAX [VGB07]. We also compare it to the dynamic power management technique presented in *Chapter 3*, to which we refer as LT-ENO. For each of the baseline implementations we use the authors' recommended parameters, i.e., $K = 3$, $D = 4$, $\alpha = 0.3$ for WCMA [PBAR09], and $\alpha = 0.5$ for EWMA [KHZS07]. For ENO-MAX [VGB07], we use $\alpha = 1/24$, $\beta = 0.25$, and $B_{target} = 65\%$. Finally, due to the hourly values given by the National Solar Radiation Database, we use $N_w = 24$ instead of 48 daily update slots for EWMA, WCMA, and ENO-MAX. This results in a slight penalty in prediction accuracy, but significantly reduces computation complexity. For the LT-ENO approach from *Chapter 3*, we assume daily control updates, i.e., $N_w = 1$, while our approach operates according to *Section 4.6.2* with weekly control updates. Moreover, note that while LT-ENO uses $W = 9$ weeks in order to trade off reactivity to variations with duty-cycle stability, the optimal algorithm's performance improves with increasing history window size, as the true average conditions are approximated with increasing accuracy. Experiments have shown that a history size of $W = 52$ weeks presents a reasonable trade-off between non-volatile memory requirements and algorithm performance.

4.7.1.2 Methodology and Input Data

To compare the proposed algorithm's performance against the baseline implementations, we leverage the simulation framework from *Chapters 2* and *3*. This framework has been shown to accurately simulate an energy harvesting system as described in [BYL⁺11], and incorporates various efficiency parameters, such as charging (η_{in}) and discharging (η_{out}) efficiencies. We let $\eta_{in} = 0.9$, and $\eta_{out} = 0.7$ respectively, but ignore all other battery inherent inefficiencies. As was already discussed in *Section 3.7* battery degradation is not considered a problem: with our approach the battery experiences a single charge/discharge cycle per year. Considering that batteries are rated for a few hundred cycles, it is expected to outlast other system components, e.g., electronics, mechanical parts, *etc.* We also enforce a low-voltage disconnect hysteresis of 60% [BAB13]. This means

that, if at any time the battery is fully depleted, the load will only be reconnected once the battery has been recharged to 60% of its capacity.

For the input data we leverage the publicly available² National Solar Radiation Database (NSRD), from where we acquired 12 years worth of trace data for each of the 8 different locations listed in *Table 4.2*. The first year for each trace is used as training data for the different approaches, while the remaining 11 years is used for simulation. Note that the data traces are given in $Wh \cdot m^{-2}$ of solar energy incident on a flat surface with zero inclination. Hence, to account for smaller panel sizes, inefficiencies of individual components, and losses in energy storage during simulation, the data is conditioned equivalently to what is described in *Chapter 2*. In *Chapter 6*, we will further exemplify our approach with data traces collected over two years with our own long-term deployment [BBF⁺11]

In order to enable a fair comparison, we provide each of the algorithms the same power subsystem, i.e., battery and solar panel. The power subsystem is obtained according to the technique described in *Chapter 2*, whereby we assume a target duty-cycle $\rho = 40\%$, and a battery capacity of $B = 100Wh$ to find the appropriate panel size from a set of four sizes, i.e., $5cm^2$, $10cm^2$, $15cm^2$, or $20cm^2$, with a panel efficiency $\eta_{PV} = 10\%$, and power level P_{max} that maximizes the resource utilization. Note that P_{max} is the maximum power dissipation at full system performance, hence a system can consume a maximum of $E_{max} = P_{max} \cdot 24$ hours of energy. This means that the utilization computed by any algorithm will be capped at E_{max} . The values used for each of the datasets are listed in *Table 4.2*. Also shown is the environmental parameter Ω , which is used for the energy approximation, and obtained from profiling the first year of each dataset.

4.7.1.3 Performance Metrics

For the performance evaluation we consider the following performance metrics:

Minimum Energy Utilization. Achieving a minimum performance level is the main objective in this study as it can be crucial in certain application scenarios, e.g., safety-critical systems. We report the minimum energy utilization for each of the approaches. Note that any outages due to a depleted battery causes the minimum to be 0, which also affects the application utility.

Application Utility. While the definition of utility is very application specific, for performance analysis we assume the utility U to be defined

²http://rredc.nrel.gov/solar/old_data/nsrdb/

Tab. 4.2: 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 for a panel with surface area $A_{pv} = 15cm^2$. P_{max} is the maximum possible system consumption, while A_{pv} is the employed panel size. Finally, Ω is the environmental parameter used for the energy estimation.

	Time Period	Lat [°]	Long [°]	Max	Avg	Min	Var	P_{max} [W]	A_{pv} [cm^2]	Ω
TX2	1/61 – 12/72	28.05	-97.39	10.08	6.47	1.75	4.37	0.4135	10	0.3431
TX1	1/61 – 12/72	31.77	-106.48	10.57	7.68	1.30	3.85	0.3915	15	0.2242
CA	1/98 – 12/09	34.05	-117.95	10.37	7.03	0.92	5.62	0.353	10	0.3282
MD	1/61 – 12/72	39.29	-76.61	10.44	5.78	0.94	6.47	0.3915	15	0.3351
MI	1/98 – 12/09	42.05	-86.05	10.55	5.34	0.53	9.05	0.286	15	0.4729
OR	1/61 – 12/72	45.52	-122.67	10.89	5.10	0.73	8.93	0.27	15	0.3955
ON	1/98 – 12/09	48.05	-87.65	10.98	5.07	0.44	11.24	0.248	20	0.4855
AK	1/61 – 12/72	61.21	-149.90	11.43	3.57	0.09	8.7	0.128	20	0.3448

as shown in Equation (4.10), where $\mu(u(\tau)) = \sqrt{u(\tau)}$.

$$U(t_1, t_2) = \sum_{\tau=t_1}^{t_2-1} \mu(u(\tau)) \quad (4.10)$$

4.7.2 Experimental Results

4.7.2.1 OPT with Energy Estimation Model

Having shown the optimality of the algorithm with the clairvoyant estimator both by example and analytically, here we investigate the performance of the proposed algorithm with the realistic energy estimator proposed in Section 4.6.1, and the datasets from Table 4.2.

Figure 4.3 shows two years of simulation results with the MI dataset, and the energy estimator discussed in Section 4.6.1. We first note the obvious clipping due to P_{max} as discussed in Section 4.7.1.2. As expected when the battery is full, the OPT algorithm follows the input (i.e., $\tilde{p}(t)$) until the net expected energy generation is less than the consumption. At that point OPT is appropriately pessimistic, while growing more optimistic in an attempt to fully leverage the battery as the winter progresses. The conservative increase in utilization is of course a desirable behavior. Nevertheless, as is visible in the figure, small negative deviation from the estimation can be fatal, particularly at times when the battery reserves are low and therefore do not suffice to make up for the mis-prediction.

By closer inspection of different datasets, we find that the estimation should be more pessimistic during critical times, i.e., when the battery is nearly empty, and note that there are two intuitive approaches to achieve this: (1) force higher utilization in fall/early winter, or (2) force lower utilization in late winter/early spring. Solution (2) makes more sense for two reasons. First it is hard to know how optimistic the system can be in the beginning of the winter, and second, for maximizing the minimum utilization, it is less critical to throttle at a point where the utilization is significantly higher than the minimum observed (which happens at the beginning of the winter). This reasoning is validated when we consider the perfect utilization achieved with a clairvoyant OPT for the same trace (see Figure 4.1). It shows that the utilization could be higher in the beginning of the winter (including summer & fall in this particular case), but should be lower at the end of winter and well into spring.

From the preceding discussion we conclude that the energy estimation should be scaled by a non-uniform scaling function $g(t)$ across different weeks, as exemplified in Figure 4.4. While likely not optimal for every possible situation, the general tendency has been empirically identified to fit well for all 8 datasets that we used for evaluation. The result of applying




Fig. 4.3: (top) Use function $u(t)$ computed by OPT for a given harvest function $p(t)$ and the estimator $\tilde{p}(t)$ from Section 4.6.1 and (bottom) the corresponding stored energy $b(t)$.




Fig. 4.4: Empirically obtained piecewise linear scaling function $g(t)$ for 1 year. Note the scales.

this scaling function on the same dataset (MI) is shown in Figure 4.5. We note that the minimum use function (ignoring the low-power outage for the moment) is reduced by only 4% from Figure 4.5 to Figure 4.4, as the minimum tends to occur in the beginning of the winter, where the scaling assumes a relatively high value.

In summary, the runtime estimation model consists of an astronomical model scaled by a dynamically adapted factor α (see Section 3.3.1), and the statically defined, non-uniform piece-wise linear function $g(t)$. With this approach we can satisfy the objective of maximizing the minimum achievable use function, while avoiding failure states.




Fig. 4.5: (top) Use functions $u(t)$ and $u_{LUT}(t)$ computed by OPT with the estimator $\tilde{p}(t)$ from Section 4.6.1 and with the LUT approach, respectively. (bottom) The corresponding stored energy $b(t)$ and $b_{LUT}(t)$.

4.7.2.2 OPT with Look-up Table

Next we discuss the behavior of the look-up-table implementation from Section 4.5. Note that we obtained \overline{F}_t as discussed in Section 4.5 using the piecewise linearly scaled estimator from the previous subsection. At runtime, the system indexes the look-up table with the current time t and battery fill-level $b(t)$ and performs a few basic operations (see Equation (4.9)) to compute the corresponding utilization. The resulting utilization $u_{LUT}(t)$ for the MI dataset is shown along-side the computationally more complex original approach in Figure 4.3. We notice that both implementations perform very similarly. However the look-up table approach tends to suffer a small penalty in minimum utilization by being too conservative early in winter. Section 4.7 further discusses the marginal performance loss due to the look-up table approach.

4.7.2.3 Comparison to LT-ENO

In this section we compare the OPT algorithm to the dynamic power management scheme from Chapter 3, to which we refer as LT-ENO. We discuss it separately from the other baseline algorithms, as it is the only algorithm that is specifically designed for solar energy harvesting systems to achieve long-term uninterrupted operation at a stable and deterministic minimum utilization. For comparison we consider the same dataset as in the previous sections. Refer to Section 4.7.2.5 for a complete discussion of the results for all datasets and algorithms.




Fig. 4.6: (top) Use functions $u(t)$ computed by LT-ENO for a given harvest function $p(t)$, and, for reference, the use function computed by the OPT approach $u_{OPT}(t)$. (bottom) The corresponding stored energy $b(t)$ for LT-ENO.

Figure 4.6 shows the use function $u(t)$ for LT-ENO, along with the input energy $p(t)$, the design-time energy approximation model $\tilde{p}(t)$. For reference, we also show the utilization with the OPT algorithm $u_{opt}(t)$.

We first note that the LT-ENO algorithm exhibits remarkably low utilization variance that is maintained throughout the year. This stems from the fact that this algorithm takes a long-term approach to compute the utilization by considering both the periods of deficit, i.e., when the generation is below consumption, and the periods of surplus, i.e., the generation exceeds consumption. This conservative approach allows the battery to bridge the periods of deficit, while simultaneously ensuring that the employed panel can actually recharge the battery during periods of surplus.

This pessimistic approach has a very negative impact on the total achievable system utility, however. From the battery profile shown on the bottom axis in Figure 4.6 it is obvious that the battery is full for most of the summer, yet the algorithm neither leverages the battery or surplus energy by increasing the utilization. This stems from a fundamentally different approach, which attempts to ensure that the battery is full at the beginning of winter, at the cost of significantly reduced total system utility. In comparison, the OPT algorithm is able to leverage the surplus energy during the summer and so maximize the total system utility, while ensuring that the battery is full at the appropriate time, and so ensure continuous operation.

Finally, we note that the minimum service-level achieved in this particular example is roughly equivalent during winter. In fact, on the period $t \in [250, 270]$ and $t \in [304, 320]$, OPT improves over LT-ENO's performance by only 0.9% and 2.8% respectively. However, when considering the entire period shown, the optimal algorithm OPT achieves roughly 12.5% higher minimum utilization than LT-ENO. This improvement is mainly due to LT-ENO's pessimistic behavior during summer, where it tends to achieve lowest utilization despite the significant energy surplus that keeps the battery at full capacity.

In summary, the comparable utilization during winter can be mainly attributed to the fact that both algorithms use a very similar estimation model, and are able to leverage the available battery capacity. LT-ENO's very pessimistic approach that fails to leverage surplus energy during summer, significantly hurts its performance. The service-level achieved by the OPT algorithm, on the other hand, is constrained only by the approximation accuracy of the estimator $\tilde{p}(t)$, and can simultaneously maximize both the minimum utilization and total system utility.

4.7.2.4 Baseline Algorithms

In the following we briefly discuss the behavior of the remaining three algorithms used for benchmarking our approach, i.e., EWMA [KHZS07], WCMA [PBAR09], and ENO-MAX [VGB07]. A detailed performance comparison is deferred to *Section 4.7.2.5*.

Recall that we parameterized each of the algorithms with the author's recommended, and arguably best parameters. Furthermore, they are given the same adequately provisioned power subsystem for simulation. However, due to limited prediction accuracy, EWMA and WCMA are required to operate at a relatively high control frequency to achieve acceptable performance. While ENO-MAX is theoretically able to run at much lower control frequencies, the original parameterization was done for half hour update periods. Since the datasets used are given in hourly values, we use $t = 1$ hour for simulation of these baseline algorithms.

EWMA. This approach uses a simple prediction scheme based on an Exponentially Weighted Moving Average (EWMA) of past observations, and an energy allocation scheme that attempts to continually correct mispredictions by increasing or decreasing the energy usage in future slots [KS03]. This algorithm is frequently used as baseline for comparison of energy prediction and allocation algorithms, and performs very well with little overhead. However, this algorithm is designed to predict energy on a horizon of a few minutes to hours, and hence will not be able to make full use of the allocated battery capacity. In fact, the algorithm may reduce

the utilization as low as ρ_{min} , despite the battery being full. Here, ρ_{min} specifies the user defined minimum acceptable utilization.

WCMA. The Weather Conditioned Moving Average (WCMA) approach has been shown to outperform the prediction accuracy of EWMA, e.g., [BBB09]. However, this does not translate into significant performance improvements with the energy allocation algorithm from [KHZS07]. This approach is significantly more computationally involved than the previous approach, but behaves very similar. Just like EWMA, it also requires a system that is capable of measuring (or somehow approximating) the energy generated by the panel.

ENO-MAX. ENO-MAX uses well-established control theory to determine the utilization that will result in maximum energy efficiency, while ensuring *Energy Neutral Operation*, i.e., consumption is always less than generation. More specifically, this algorithm adjusts the utilization such that a pre-defined, static battery threshold B_{target} may be maintained. While it is very effective in achieving its maximization objective, the *Energy Neutral Operation* objective causes the algorithm to aggressively adjust the utilization when B_{target} may not be maintained and therefore leads to high variability. On the one hand, the ability to maximize the energy usage when the battery fill level $b(t)$ is above B_{target} , leads to high overall utility. On the other hand, the aggressive downscaling when $b(t) < B_{target}$ results in minimum utilization that follow a user-specified minimum accepted level ρ_{min} .

Clearly, the ability to satisfy ρ_{min} depends on proper selection of B_{target} and ultimately on the proper battery and panel size. In fact, we believe that the main weakness of this approach is the use of a static threshold value B_{target} . The algorithm could perform much better with a dynamically computed setpoint, possibly based on the same energy model that is used in this chapter.

4.7.2.5 Performance Comparison

In this section we discuss the performance results for the proposed approach, as well as the previously introduced algorithms for all datasets listed in Table 4.2. The results are summarized in Table 4.3. Note that for EWMA, WCMA, and ENO-MAX, we report the result for both a pessimistic $\rho_{min} = 1\%$ (as is done in e.g., [VGB07, KHZS07]), and, listed in parentheses in Table 4.3, a more realistic value of $\rho_{min} = 30\%$ (recall that the power subsystem is designed with an expected duty-cycle $\rho = 40\%$).

First, as expected the EWMA and WCMA algorithms perform comparably. This is not surprising, as the same energy allocation scheme is used. Hence the difference stems from differing prediction accuracies. In fact, WCMA achieves a better minimum utilization only for two datasets

(ON, AK), but improves over EWMA by up to 20.5% in system utility for $\rho_{min} = 1\%$ and 6.5% for $\rho_{min} = 30\%$. Both of these algorithms are unable to leverage the battery, causing the minimum achieved utilization to follow the user-defined ρ_{min} , even when the battery is full. However, as has been shown in *Chapter 3*, if the battery is not appropriately provisioned, achieving even ρ_{min} may not be possible.

The ENO-MAX algorithm has a similar limitation. While achieving the highest utility of all algorithms, the minimum utilization varies greatly, and tends to be the lowest of all evaluated algorithms when using $\rho_{min} = 1\%$. For $\rho_{min} = 30\%$ the minimum performance is equivalent to that achieved of EWMA and WCMA. This is to be expected, as the range of acceptable utilization to chose from is significantly reduced. In fact, the utilization computed by these algorithms is overridden by ρ_{min} most of the time. The pessimistic behavior in terms of minimum utilization of these three algorithms is attributed to their battery agnostic nature.

Next, the LT-ENO achieves significantly higher minimum utilization than the algorithms discussed so far. In the best case, LT-ENO outperforms ENO-MAX (with $\rho_{min} = 1\%$) by 3660%. In the worst case, the minimum service-level achieved by LT-ENO is only approximately 0.7% better than that of the baselines with $\rho_{min} = 30\%$. While the achieved utility is generally higher than EWMA and WCMA, it lags behind ENO-MAX. At worst, the utility achieved by LT-ENO is 32% below ENO-MAX, and at best up to 64% higher than EWMA.

The OPT algorithm with the explicit energy model outperforms all approaches in terms of minimum service-level, except with the TX2 dataset, for which it performs only 4% worse than LT-ENO, but still better than all other baselines. In fact, our approach comes to within 9.9% of the clairvoyant approach (CV) at best, and 29.5% at worst. The utility achieved by OPT is only surpassed by ENO-MAX, which exhibits significantly lower minimum utilization and significant variability.

Finally, the look-up table implementation performs negligibly worse than the approach that explicitly computes the algorithm. This demonstrates that our approach can offload the computational effort to design-time, and achieve very good performance at the cost of a few simple operations once a week with the values in the look-up table.

4.7.3 Practical Considerations: Energy Estimation

The computational overhead of our solution has been detailed in *Section 4.5*. Here we briefly discuss practical considerations related to the energy estimation. *Chapter 3* already presented considerations regarding measurement support, global time knowledge, and battery inefficiencies.

Tab. 4.3: Summary of performance results for the clairvoyant optimal algorithm (CV), the optimal algorithm with the energy estimator from *Section 4.7.2.1* (OPT), the look-up table implementation (LUT), and the baselines LT-ENO, ENO-MAX, EWMA, and WCMA. Values in parentheses are for $\rho_{min} = 30\%$. For EWMA, WCMA, and ENO-MAX, the absolute minimum utilization follows the user-defined minimum acceptable performance level ρ_{min} .

		Dataset							
		TX2	TX1	CA	MD	MI	OR	ON	AK
CV	Min	3.62	4.19	3.76	3.88	3.04	2.63	2.76	1.58
	U	3235.2	3625.4	3439.4	3603.1	3453.6	3349.5	3612.3	3070.9
OPT	Min	2.86	3.76	3.42	3.38	2.73	2.36	2.32	1.22
	U	3026.1	3399.9	3244.1	3407.5	3272.0	3139.6	3210.0	2336.4
LUT	Min	3.01	3.71	3.01	3.03	2.53	2.35	2.30	1.21
	U	3023.7	3474.1	3177.0	3437.7	3278.2	3086.4	3201.7	2320.9
LT-ENO	Min	3.00	3.66	2.98	3.06	2.4	2.04	2.19	1.09
	U	2854.5	3082.6	2857.8	2806.5	2549.1	2380.5	2385.8	1727.4
ENO-MAX	Min	0.10 (2.98)	0.10 (3.13)	0.54 (2.54)	0.09 (2.82)	0.52 (2.06)	0.06 (1.93)	0.66 (1.79)	0.03 (0.92)
	U	3198.4 (3205.4)	3600.3 (3601.2)	3409.1 (3410.3)	3566.7 (3568.5)	3317.9 (3317.7)	3143.5 (3147.1)	3170.1 (3167.6)	2236.5 (2248.4)
WCMA	Min	0.18 (2.98)	0.19 (3.13)	0.15 (2.54)	0.17 (2.82)	0.12 (2.06)	0.12 (1.93)	0.11 (1.79)	0.06 (0.92)
	U	1989.0 (2878.7)	2266.8 (3073.2)	2104.1 (2815.9)	2101.5 (2796.6)	1858.0 (2505.2)	1789.4 (2322.1)	1706.2 (2312.6)	1356.4 (1712.2)
EWMA	Min	0.18 (2.98)	0.19 (3.13)	0.15 (2.54)	0.17 (2.82)	0.12 (2.06)	0.12 (1.93)	0.06 (1.79)	0.03 (0.92)
	U	1735.2 (2766.0)	1877.0 (2887.0)	1859.1 (2686.2)	1907.4 (2716.2)	1775.8 (2490.9)	1616.7 (2313.8)	1549.2 (2291.8)	1124.9 (1652.7)

It is clear that the performance achievable by the algorithm is ultimately limited by the estimation. Finding an appropriate energy estimation model that is neither too conservative nor too optimistic has been shown to be difficult, as the weather patterns that affect the solar energy incident are hard to model [Buz08] and difficult to predict [HLG06]. Nevertheless, the non-uniformly scaled energy model proposed in *Section 4.7.2.1* allows the algorithm to achieve a minimum performance level to within 9.9% of the optimal (with clairvoyant estimator) at best, and 29.5% at worst. Similarly, the application utility achieved with our ap-

proach reaches the optimal utility to within 5.5% at best, and 31.4% at worst. While the estimation model, and particularly the scaling function may be fine tuned for the different datasets to improve these results, we opted to use one model for all datasets.

4.8 Chapter Summary

We have shown that the proposed approach successfully achieves the objective of providing a guaranteed minimum energy utilization, and therefore minimum service-level. In fact, we believe to be the first to provide an analytic solution for solar energy harvesting embedded systems, which ensures uninterrupted operation at a maximized minimum utilization. Using eight datasets, we show that our approach significantly outperforms four baseline implementations both in terms of minimum utilization and total utility. In *Chapter 6* we will provide further performance results and show that the proposed scheme is highly effective for realistic scenarios. Finally, it is important to note that our algorithm is applicable to any harvesting source, presuming that an adequate energy estimator can be defined.

5

State-of-Charge Approximation

5.1 Introduction

In the previous two chapters we presented novel runtime energy management schemes with the objective of maximizing the minimum system utility for solar energy harvesting systems. For these schemes to work, the system requires a means to approximate the harvested energy *and* residual charge stored in the battery at any given instant in time. In this chapter we therefore present an approach to approximating the battery's fill-level, which, as we will show in *Chapter 6*, can also be used to infer the harvested energy without relying on dedicated, special-purpose hardware.

It has been shown that awareness of the energy available to the individual motes [PLR05], and the entire network [CD05] can significantly improve overall system lifetime and utility. However, accurately determining the battery fill level, referred to as battery *State-of-Charge*, presents a non-trivial problem to solve. This is because a battery's State-of-Charge depends on many battery internal and external factors, such as size and type of battery, the rate at which it is being discharged, as well as temperature and battery condition (e.g., age, present State-of-Charge, *etc.*). Nevertheless, the implementation cost of a State-of-Charge approximation technique is likely outweighed by the possible improvements in system observability, predictability, and utility [PAG09, TJC08], which stem from runtime knowledge of the battery's State-of-Charge.

Many existing approaches to determine a battery's State-of-Charge depend on dedicated hardware to monitor the energy flux into and out of the

battery. However, addition of *special-purpose hardware* not only increases design complexity but also system cost [DOTH07]. Furthermore, upgrading existing systems with the hardware necessary may not be practical or even possible. Other approaches use *complex software models* that incorporate the non-linear characteristics of batteries under varying and uncontrollable conditions. However, many of the proposed models present significant configuration and computational overhead [RVR03]. Considering the limited computational resources available on typical wireless sensor network platforms, these models may thus not be feasible for implementation on such systems.

Contributions

In this chapter we present a practical trace-based, direct-measurement [BKN02] method for online battery State-of-Charge approximation. It is aimed at off-the-shelf photo-voltaic harvesting setups, and therefore does not rely on special purpose hardware, but only requires low-cost sensors that are commonly available on contemporary wireless sensor network platforms for system health monitoring. Moreover, the method is computationally inexpensive as it does not depend on complex battery models. Instead, we leverage the known characteristics of the charge controller and the battery's behavior under load.

More specifically, the contributions of this chapter can be summarized as follows.

- We present a practical, trace-based, direct-measurement [BKN02] method for online battery State-of-Charge approximation that does not rely on special purpose hardware.
- We show how leveraging known characteristics of the battery behavior under load can be used to devise a computationally light-weight method to State-of-Charge approximation.
- We present a state machine for tracking the different phases of the charging and discharging process at runtime.
- We show how the battery can be profiled such that the model parameters can be defined even for batteries that differ from the profiled battery in terms of battery age and capacity, as well as operating temperature.
- We perform extensive experiments under different conditions, such as operating temperature, load variations, and battery age.

- With the experiments we show that our direct-measurement approach achieves State-of-Charge approximation with an average error below 5% when compared to the State-of-Charge inferred by the measurements of a coulomb counter.
- Finally, we demonstrate that accurate lifetime predictions even with temperature fluctuations and varying battery, and load conditions are possible, as our method implicitly accounts for battery inefficiencies such as temperature and aging.

Roadmap

The remainder of this chapter is structured as follows. In *Section 5.2* we discuss common approaches to battery State-of-Charge determination, and adaptations to wireless sensor network scenarios. In *Section 5.3* we provide a conceptual overview of the approach, while the system architecture and battery characteristics that form the foundation of the proposed approach are discussed in *Section 5.4*. In *Section 5.6* we then discuss the runtime State-of-Charge approximation in detail. *Section 5.7* describes the experimental set up and presents a performance evaluation of the approach. In *Section 5.8* we elaborate on runtime measurement considerations. Finally, *Section 5.9* presents the concluding remarks.

5.2 State-of-Charge Approximation Methods

Efforts to estimate the residual charge contained in a battery have been made for almost as long as rechargeable batteries exist, and increasing popularity of hybrid electric vehicles has further pushed this research topic. Hence, literature review reveals a host of battery State-of-Charge approximation methods, of which common techniques are presented in this section. A brief review of wireless sensor network specific approaches is also provided.

General State-of-Charge approximation methods

One of the earliest techniques is in the form of current integration, usually referred to as coulomb counting [PBNR05, PPJ01, BKN02]. With this approach, charge flowing from (in case of discharging), or to the battery (charging) is measured with a dedicated piece of hardware, and integrated over time. After subtracting the net charge from the capacity of a full battery, the residual charge contained in the battery can be obtained. Although widely used today, this approach suffers from a number of issues. First, to yield an acceptable accuracy, battery inefficiencies that are

not directly measurable must be compensated for. Second, inaccurate current readings lead to an accumulation of error. However, accurate current measurements are expensive, and the addition of the necessary hardware can add substantial design complexity and development cost. Finally, since the performance depends on known initial conditions, regular re-calibrations are necessary [PPJ01, PBNR05].

A technique that attempts to improve the performance of coulomb counting is called book-keeping [BKN02]. With this approach, common battery inefficiencies, such as discharging efficiency, self-discharge, and capacity losses are taken into consideration through software, and so yield a more accurate State-of-Charge indication. The battery inefficiencies, and details on how the model proposed in this chapter copes with their effects are discussed in *Section 5.7*.

A discharge test under controlled conditions is considered to be the most reliable means of approximating a battery's State-of-Charge [PPJ01], and commonly performed by the battery manufacturer [Con11]. For this reason, discharge tests with a coulomb counter are used as ground-truth for the evaluation of the proposed approximation method. Nevertheless, discharge test during the deployment are not feasible for most if not all practical applications because system operation must be interrupted willingly and possibly for long periods of time.

Another technique that is applicable to vented lead-acid (LA) batteries with liquid electrolytes is in the form of measuring the electrolytes' physical properties. LA batteries, which are common in cars and stationary photo-voltaic energy harvesting systems, exhibit a relationship between the electrolyte's properties and the battery's State of-Charge. Therefore, by measuring e.g., specific gravity, conductivity, and ion-concentration, the battery's State-of-Charge can be determined. However, to obtain an accurate approximation, the temperature must be considered, and measurements may only be taken after a proper charging cycle and appropriate resting period of the battery. This, together with the need for measuring the battery's internal properties limit the applicability of this approach.

A lead-acid battery's electrolytes' properties are directly related to the open-circuit voltage [Con11]. Hence, the State-of-Charge can be obtained by measuring the battery's terminal voltage when no load is connected. However, this necessitates periodic load disconnects, which requires special design considerations. Just like with the previous approach, appropriate resting times are required to obtain meaningful measurements.

Recently, there has been significant efforts in devising analytic models for batteries and State-of-Charge approximation. An excellent review is given in [RVR03]. These models range from electric circuit

equivalents, over stochastic load modeling, to mixed models that incorporate both experimental data and physical laws. In particular, extended [VPB07], and adaptive [HKS09] Kalman filtering techniques have emerged. Unfortunately, these approaches require significant configuration efforts and exhibit computational complexity and therefore preclude application in low-performance systems such as wireless sensor network motes [RVR03].

Finally, the approach most closely related to the one discussed in this work consists of direct voltage measurements [PBNR05]. Voltage measurements are generally considered to yield inaccurate State-of-Charge indication because a battery's voltage profile depends on discharge rate, temperature effects, and age of the battery to name a few. However, as will be elaborated in the remainder of this work, if the battery's behavior under load, and the discharge current are known (or measurable), the State-of-Charge can be obtained with relatively good accuracy even under varying battery operating conditions.

Wireless sensor network specific State-of-Charge approximation. Of the approaches discussed so far, only a subset is feasible for implementation in wireless sensor network scenarios. This is mainly due to cost considerations, physical constraints, and limited processing power available on the motes. In addition to techniques that leverage special-purpose hardware [RKH⁺05, JDCS07], fully software [CPBA12, KPA⁺10, DOTH07] based approaches to State-of-Charge approximation have also been proposed, and are briefly introduced in the following.

In [CPBA12] a software based model for lifetime prediction is presented in the context of solar energy harvesting. The platform load is characterized by a constant average current for each activity (i.e., sensing, processing, and transmission). The State-of-Charge is then approximated in software by "counting" the charge over the period of time a given component is active. The authors report a lifetime estimation error below 10%. However, their approach requires measurements of the current generated by the energy solar panel, or, alternatively a light meter to approximate the current. Both require pre-deployment design considerations, which precludes its use in existing systems.

Although not aimed at energy harvesting systems, [DOTH07] and [KPA⁺10] follow a similar approach. Software routines are implemented that are executed every time a certain hardware component is switched on or off. The total time a given component is active is then multiplied by its average current drain, which is obtained by pre-deployment power profiling of the system. Finally, to obtain the overall system energy consumption (and hence infer the State-of-Charge), all components are summed up. It is not clear what accuracy either of these approaches

achieve, but considering that coulomb counting Integrated Circuits require frequent re-calibration, it can be assumed that the error in purely software counting techniques accumulates just as rapidly.

In contrast to purely software based approaches, [RKH⁺05] presents Heliomote, a custom energy harvesting hardware module as add-on for Berkley/Crossbow motes. It is a complete power management solution that autonomously controls solar harvesting, charging of the battery, storage and power routing, and provides harvesting and battery state information to the host platform. While not exclusively aimed at providing State-of-Charge information, the authors show that harvesting aware power management can significantly improve system utility.

Similarly, SPOT [JDCS07] is another application specific custom micro power meter for energy monitoring of the popular MicaZ motes. As motivation, the authors state that commercially available Integrated Circuits are not designed to meet WSN application requirements. Similar to Heliomote, the immediate aim is not battery State-of Charge approximation, but rather empirical on-site evaluation of low power designs at scale. A price tag of \$25 is stated, but it is not clear how much re-design and integration effort would be necessary to adapt the SPOT module to other mote platforms.

5.3 Concept

When a battery is discharged, its terminals exhibit a voltage profile over time that is characteristic to the specific battery chemistry, the magnitude of the discharge current, the condition of the battery, and operating temperature. In this chapter we leverage this fact and devise a battery State-of-Charge approximation model that is based on low-cost voltage and current sensor readings. Our approach consists of an offline, pre-deployment phase to profile the battery, and an online, runtime phase during which the instantaneous battery State-of-Charge is computed.

The offline phase consists of a three-step procedure: we first profile the battery with different discharge (charge) currents (*Section 5.5.5.1*) to obtain battery terminal voltage and discharge (charge) current traces. As a second step (*Section 5.5.5.2*) we transform the traces from voltage as a function of time to Depth-of-Discharge as a function of voltage. For this we take the charge controller's characteristics into consideration to define 100% Depth-of-Discharge (equivalent to 0% State-of-Charge, since Depth-of-Discharge is the complement of State-of-Charge). Finally, we extract the required parameters for runtime discharge (charge) approximation from the recorded traces (*Section 5.5.5.3*).




Fig. 5.1: High-level system architecture with off-the-shelf energy harvesting components (solar panel, battery, charge controller), and electrical load, e.g., a mote. I_{sys} and V_{sys} are assumed to be observable. Note the charge controller's low-voltage disconnect switch.

At runtime, the system periodically obtains system input voltage and current measurements and executes a Finite-State Machine (see *Section 5.6.1*) to track the battery's states (discharging, bulk charging, *etc.*). Then, depending on the battery's current state, the instantaneous State-of-Charge is computed with the respective equation developed in *Sections 5.5.3* and *5.5.4* and the coefficients obtained during the offline phase.

5.4 System Architecture and Assumptions

For the proof-of-concept implementation discussed in this chapter, we assume an off-the-shelf energy harvesting set-up illustrated in *Figure 5.1*. It consists of an energy harvesting module, such as a solar panel e.g., [cle], a low-cost Pulse-Width Modulation charge controller, e.g., [Mor09], and one or multiple Valve Regulated-Absorbent Glass Mat (VR-AGM) sealed lead-acid batteries, e.g., [Con11]. While this is not a typical set-up for low-power wireless sensor network motes with simple sense-and-transmit applications it can be considered a reasonable, low-cost set-up for systems with high-energy consumption due to, e.g., high-power sensors and/or increased duty-cycles. Aside from wireless sensor network basestations [MOH04] that usually have considerable energy demands on the order of Watts, recently proposed application scenarios like continuous GPS [BSB12], acoustic emission [WGGB12], or Audio/Video surveillance [AAP⁺12] further validate such a set-up.

With such an off-the-shelf set-up, the charge controller regulates proper charging of the battery, but does not provide State-of-Charge information. Therefore, the proposed approach aims at approximating the State-of-Charge in a battery type, and set-up independent manner. The goal is to provide a solution that does not depend on extensive hardware support, but gets by with low-cost sensors commonly available




Fig. 5.2: Manufacturer provided Depth-of-Discharge versus terminal voltage [Con11] for 1, 2, 4, 8, 20, and 120-hour discharge rate (as a percentage of 20-hour rate) of a VR-AGM battery.

on contemporary wireless sensor network motes for health monitoring purposes. Furthermore, accurate State-of-Charge indication should be possible without requiring extensive processing at runtime.

Although we assume a lead-acid battery chemistry in this work, the approach applies to other battery chemistry with an appropriate discharge profile, i.e., similar to what is shown in *Figure 5.2*. In [Ste13] we validated the approach discussed herein with Nickel-Metal Hydride (NiMH) batteries and a TinyNode [DFFMM06] wireless sensor network mote depicted in *Figure 1.1b*, and show that for NiMH batteries minor modifications to our model are necessary. Moreover, the proposed approach is agnostic to a particular set-up and harvesting source, and only requires a well-defined behavior of the charge-controller and battery under load.

5.5 Battery Model

Devising an analytic battery model that takes all the battery's non-linearities into consideration has proven to be complex [BKN02] and, depending on application, computationally prohibitively expensive [RVR03]. However, incorporating all inefficiencies into the model

may not be necessary for achieving acceptable State-of-Charge approximation. Instead, from the illustration of a battery's voltage profile (obtained from its data-sheet [Con11]) in *Figure 5.2*, we observe that this particular battery type exhibits qualitatively very similar voltage versus Depth-of-Discharge¹ curves for different discharge rates. In the following we describe how this observation can be leveraged to approximate the battery's State-of-Charge by observing only the battery's terminal voltage and system drain current.

In *Section 5.5.1* we first define a battery's rated capacity and the charge controller's role in our model. In *Section 5.5.2* we present a high-level discussion of the battery's charging and discharging characteristics. Then, in *Sections 5.5.3* and *5.5.4* we explain how these characteristics are leveraged to devise a battery model for the charging and discharging process. Finally, in *Section 5.6* we present the runtime battery State-of-Charge approximation approach.

5.5.1 Capacity and Cut-off Voltage

The manufacturer provided capacity rating specifies the battery's nominal capacity C , which refers to the maximum charge that can be withdrawn before fully depleting the battery (i.e., 100% Depth-of-Discharge) at a specified discharge rate and temperature. Usually, the discharge rate and temperature are assumed to be $C/20$ and 25°C respectively, which means that the battery is discharged at a rate such that it is fully depleted in 20 hours at a temperature of 25°C . However, since deep discharge cycles cause irreversible chemical reactions, the battery should not be fully discharged. Therefore, in addition to regulating proper, i.e., safe, charging of the battery, commercial charge controllers implement a cut-off voltage V_{cutoff} , below which the load is disconnected until the battery recovers to a certain State-of-Charge (usually 60% of nominal capacity). This mechanism protects the battery from deep-discharge cycles and thus maximize its lifetime.

The known cut-off voltage provides two pieces of information that we leverage in the proposed model. First, the cut-off voltage is used to define 0% State-of-Charge (equivalent to 100% Depth-of-Discharge) of the battery. Second, since the assumed charge controller sets this voltage such that the Depth-of-Discharge is limited to 80%, if we ignore the battery's discharging efficiency, it can be assumed that the battery can deliver 80% of the nominal capacity. This assumption has been experimentally verified with a coulomb counter [Max98] in *Section 5.7*.

¹Depth-of-Discharge DoD is the complement of State-of-Charge SoC , i.e., $SoC = 1 - DoD$




Fig. 5.3: Qualitative illustration of the relation between battery terminal voltage (top) and current (bottom) during the charging and discharging processes.

5.5.2 Charge and Discharge Profiles

The behavior of a battery can be classified into two processes: *charging* and *discharging*. The charging process can be further segmented into three phases, enforced by the employed charge controller (SunSaver SS-6L): *Bulk*, *Absorption*, and *Float* charging. These processes are illustrated in Figure 5.3 and briefly introduced in the following. The battery and charge controller's characteristics during these phases will be exploited in Sections 5.5.3 and 5.5.4 where we devise a trace-based model to approximate the battery's State-of-Charge based on voltage and current measurements. For simplicity, we assume ideal conditions, i.e., no charging inefficiencies, and a stable power source.

Discharging. The battery is discharged with a certain drain current, e.g., $I_{\text{drain}} = C/20$, that defines the rate of discharge. Electro-chemical reactions lead to an increase in battery internal resistance, which causes a drop of the battery's terminal voltage directly related to the magnitude of the drain current. As illustrated in Figure 5.3, if the battery experiences a load change, i.e., the drain current changes in magnitude, the battery voltage reacts accordingly. In case of an increase in drain current, the battery voltage drops. Similarly, if the load is reduced, the drain current decreases and the battery's terminal voltage raises accordingly. It is important to note that the load caused by the

drain current $I_{\text{drain}}(t)$ is relative to the battery's nominal capacity. For example, a high current on a high-capacity battery imposes a similar load as a proportionally lower current on a proportionally lower capacity battery. For this reason, we introduce the relative load $RL(t)$, which normalizes the discharge current with the respect to the battery's nominal capacity, i.e., $RL(t) = I_{\text{drain}}(t)/C$. As we will show in *Section 5.5.3*, this allows us to devise a State-of-Charge approximation model that is independent of the battery's capacity.

Bulk Charging. In this phase, the charge controller enforces a constant current I_{bulk} , which causes the battery potential to increase until a pre-defined, temperature dependent voltage V_{bulk} is reached. The time for reaching this voltage depends on the magnitude of the net charging current and the battery's State-of-Charge at the beginning of the bulk charging phase.

Absorption Charging. In this phase the charge controller keeps the battery terminals at a constant voltage level V_{abs} defined by the bulk phase. The charging current $I_{\text{abs}}(t)$ is tapered off exponentially until it drops below a certain percentage of the battery's nominal capacity (0.5% for the batteries assumed in this chapter [Con11]), after which the battery can be considered fully charged.


Float Charging. During the final charging phase, the charge controller attempts to keep the battery at 100% State-of-Charge. To prevent over-charge, the voltage is reduced to a temperature dependent, constant float voltage V_{float} and the charging current I_{float} is kept constant. Depending on the State-of-Charge at the beginning of the charge process, and the magnitude of the charging and drain currents, the absorption and float charging phases may not be reached.

5.5.3 Battery Discharging Model


In this section we devise a model that can approximate the State-of-Charge given measurements of the system input voltage and drain current. To this end we develop a set of equations to model the discharging process based on the observations from *Section 5.5.2*. In *Section 5.5.4* we present the same for the charging process, while *Section 5.6* discusses how the system can determine the different phases of the charging/discharging process at runtime.

Figure 5.4a provides a qualitative illustration of the battery's voltage profiles over time when discharged with two different relative loads. As discussed in the previous section, the voltage observed is related to the magnitude of the relative load $RL(t)$, and, as illustrated, the time until the charge controller performs a low-voltage disconnect is lesser with a higher load. The start of each trace is indicated with point *A*, at which the

battery is assumed to be fully charged. The time when the battery reaches the cut-off voltage is denoted t_{cutoff} , at which point the battery can be considered empty (points B and C).



(a) Discharge traces at two different load levels. The charge controller specific cut-off voltage is indicated by points B and C.



(b) Voltage vs. *DoD* curves obtained through domain transformation and inversion of the discharge traces shown in Figure 5.4a.

Fig. 5.4: (a) Qualitative illustration of discharge traces, and (b) *Depth-of-Discharge* approximation as a function of voltage. Labels A–D are given to illustrate the transformation from Figure 5.4a to Figure 5.4b

In the following we assume that i such discharge profiles for the particular battery chemistry and relative loads $RL_i(t)$ are available. In Sec-

tion 5.5.5 we will discuss how to appropriately profile the battery such that the proposed model can be parameterized adequately.

As is evident from *Figure 5.4*, the profiles for the two different loads are given in terms of battery voltage over time. Since we would like to obtain a representation of State-of-Charge as a function of battery voltage and drain current, we perform the following procedure.

First, as illustrated in *Figure 5.4*, we transform the traces from voltage as a function of time (*Figure 5.4a*) to Depth-of-Discharge as a function of voltage (*Figure 5.4b*). To this end we use the known cut-off voltage of the charge controller to define 100% Depth-of-Discharge, and the known battery voltage of a fully charged battery to define 0% Depth-of-Discharge. Further, since we assume the traces to be obtained with a constant relative load (see *Section 5.5.5*), we use the recorded discharge time $t_{cutoff,i}$ until the battery reaches the cut-off voltage to assign the Depth-of-Discharge linearly in time, i.e., $DoD_i(t) = t_i/t_{cutoff,i}$. This effectively transforms the traces from the time domain to the Depth-of-Discharge domain, which is defined for the interval $[0,1]$. Finally, we inverted the traces to obtain the Depth-of-Discharge as a function of the battery voltage, as illustrated in *Figure 5.4b*.

After transformation of the traces, *Figure 5.4b* suggests that we should be able to closely approximate the battery voltage dependent Depth-of-Discharge with a polynomial function of order n . The choice of n is a trade-off between computational complexity and profiling overhead versus runtime accuracy of the model, but it has been found that a quadratic approximation achieves very good results, except when the battery is nearly full (see *Section 5.7*). Nevertheless, as is intuitively clear, each traces' polynomial approximation, shown in *Equation (5.1)*, allows for Depth-of-Discharge estimation *only* for the relative load with which the trace was recorded.

$$\widehat{DoD}_i(V_{bat,i}) = a_{n,i} \cdot V_{bat,i}^n + a_{n-1,i} \cdot V_{bat,i}^{n-1} + \dots a_{1,i} \cdot V_{bat,i} + a_{0,i} \quad (5.1)$$

However, we wish to represent the Depth-of-Discharge with a single equation that covers the entire operating range, including relative loads other than the ones used to record the traces. For this we leverage the fact that the polynomials' coefficients $a_{n,i}$ are dependent on the respective relative load RL , and rewrite *Equation (5.1)* as shown in *Equation (5.2)*.

$$\widehat{DoD}(V_{bat}, RL) = a_n(RL) \cdot V_{bat}^n + a_{n-1}(RL) \cdot V_{bat}^{n-1} + \dots a_1(RL) \cdot V_{bat} + a_0(RL) \quad (5.2)$$

Then, to obtain the coefficient functions $a_n(RL)$, we fit $n+1$ polynomial approximations of order $m \geq n$ over all i coefficients from *Equation (5.1)* to find the coefficients $b_{n,m}$ in *Equation (5.3)*. Clearly, the approximation order m must be chosen such that the resulting interpolation best fits

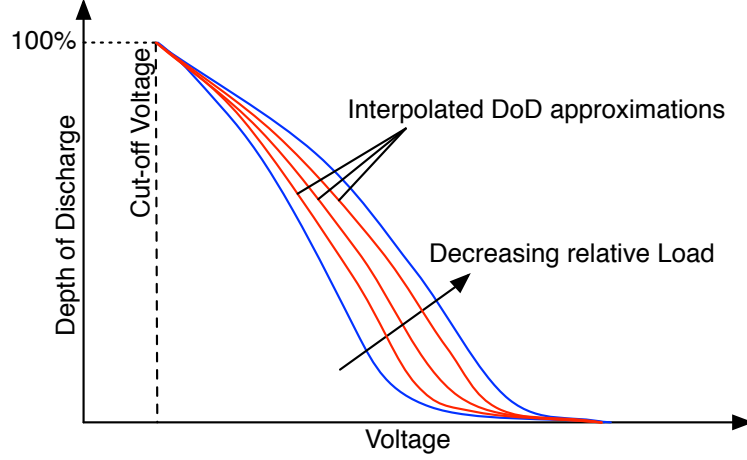


Fig. 5.5: Interpolated Depth-of-Discharge approximations by relative load.

all coefficients, which implies that at least $m + 1$ traces are recorded (see Section 5.5.5).

$$a_n(RL) = b_{n,m} \cdot RL^m + b_{n,m-1} \cdot RL^{m-1} + \dots + b_{n,1} \cdot RL + b_{n,0} \quad m \geq n \quad (5.3)$$

A qualitative illustration of the interpolated polynomial Depth-of-Discharge approximation is shown in Figure 5.5. The outermost two curves (shown in blue) approximate the recorded traces, while the remaining three curves (shown in red) are approximations with interpolated coefficients for relative loads other than those used for recording the traces, i.e., RL_i . As we will discuss in Section 5.5.5, the recorded traces should bound the expected load of the system, i.e., $RL_{sys,max} \geq RL_i \geq RL_{sys,min}$.

In summary, with the approach discussed so far, we characterize the battery discharging process with a single expression that represents the Depth-of-Discharge as a function of relative load and battery voltage, and $n + 1$ equations for the coefficient functions $a_n(RL)$. As we will explain in Section 5.6, to obtain the instantaneous Depth-of-Discharge approximation $\widehat{DOD}(\cdot)$ at runtime, the system evaluates Equations (5.2) and (5.3). Based on the result, the State-of-Charge approximation $\widehat{SOC}(\cdot)$ can be trivially obtained as $\widehat{SOC}(\cdot) = 1 - \widehat{DOD}(\cdot)$.

5.5.4 Battery Charging Model

Just as for the the *discharging* process, we want to approximate the State-of-Charge during the *charging* process with current and voltage measurements. However, since we assume an off-the-shelf set-up, the charging current $I_{PV}(t)$ produced by the solar panel cannot be measured directly

and must therefore be approximated. Moreover, depending on the magnitude of $I_{PV}(t)$, and assuming non-zero drain current $I_{sys}(t)$, the net current flowing into the battery may be positive (i.e., battery is charging, $I_{sys}(t) < I_{PV}(t)$), negative (i.e., battery is discharging, $I_{sys}(t) > I_{PV}(t)$), or zero, i.e., $I_{sys}(t) = I_{PV}(t)$. In the case of $I_{sys}(t) > I_{PV}(t)$, the model from the previous section applies. In the following we discuss how we model the charging process without requiring explicit knowledge of the charging current.

5.5.4.1 Bulk Charging Model

As illustrated in *Figure 5.3*, in the ideal case, the bulk charging current, $I_{bulk} \leq I_{PV}$ flowing into the battery is constant, causing the voltage to increase monotonically. Just as in the previous section, we use the relative charging current, i.e., $RC_{bulk} = (I_{PV} - I_{sys})/C$, to determine the current State-of-Charge during the bulk phase in a conservative way (conservative because the current is relative to nominal, rather than apparent capacity, which is only a fraction of the nominal capacity, see *Section 5.5.1*). However, the exact value for I_{PV} is unknown, and RC_{bulk} must therefore be approximated. As illustrated in *Figure 5.6*, which shows qualitative voltage traces during the bulk charging phase for two different charging currents $I_{bulk,1}$ and $I_{bulk,2}$, where $I_{bulk,1} < I_{bulk,2}$, the rate at which the voltage increases is directly proportional to the net bulk charging current. Therefore, we can estimate the relative charge with a first order approximation as shown in *Equation (5.4)*. The slope $\delta V/\delta t$ of the voltage profile can be obtained with a linear regression. The parameterization of *Equation (5.4)* will be discussed in *Section 5.5.5*.

$$\widehat{RC}_{bulk} = c_1 \cdot \frac{\delta V}{\delta t} + c_0 \quad (5.4)$$

The change in State-of-Charge due to \widehat{RC}_{bulk} is relative to the State-of-Charge before entering the *bulk* phase, i.e., the last State-of-Charge approximation during the *discharge* phase. Thus, the instantaneous State-of-Charge is computed as shown in *Equation (5.5)*, where t_{bulk} represents the time elapsed since entering the bulk phase.

$$\widehat{SoC}(t) = \widehat{SoC}(t - t_{bulk}) + \widehat{RC}_{bulk} \cdot t_{bulk} \quad (5.5)$$

5.5.4.2 Absorption Charging Model

Due to the constant terminal voltage during the absorption phase (see *Figure 5.3*), the relative charge $RC_{abs}(t) = (I_{PV}(t) - I_{sys}(t))/C$ cannot be




Fig. 5.6: Linear approximation of the voltage slope of two different net bulk charging currents.

estimated through the battery's terminal voltage in the same way it is done for bulk charging. However, as explained in *Section 5.5.2*, and assuming an ideal scenario, i.e., the current source does not fluctuate significantly, the absorption phase is characterized by an exponentially decreasing relative charge with initial value given by \widehat{RC}_{bulk} . We leverage this fact, and estimate the true relative charge $RC_{abs}(t)$ during absorption as shown in *Equation (5.6)*, where $\widehat{\lambda}$ is the decay constant, which is dependent on the relative charging current during the bulk phase.

$$\widehat{RC}_{abs}(t) = \widehat{RC}_{bulk} \cdot e^{\widehat{\lambda}(\widehat{RC}_{bulk}) \cdot t} \quad (5.6)$$

Since \widehat{RC}_{bulk} has been defined in the previous section, we only need to define $\widehat{\lambda}(\cdot)$. For this we first find the true decay constants $\lambda_i(RC_i)$ for each of the i charging traces with an exponential fit over the corresponding trace. Then we can define $\widehat{\lambda}(\cdot)$ with a linear approximation over the i true decay constants to obtain the expression shown in *Equation (5.7)*.

$$\widehat{\lambda}(\widehat{RC}_{bulk}) = d_1 \cdot \widehat{RC}_{bulk} + d_0 \quad (5.7)$$

Finally, we can determine the State-of-Charge in a similar manner as was done for the bulk phase. However, the relative charge during the *absorption* phase is decreasing over time, hence we compute the State-of-Charge differentially for each discrete approximation time step Δt , as shown in *Equation (5.8)*.

$$\widehat{SoC}(t) = \widehat{SoC}(t - \Delta t) + \widehat{RC}_{abs}(t) \cdot \Delta t \quad (5.8)$$

5.5.4.3 Float Charging Model

The State-of-Charge during the float phase is by definition 100%. As discussed in *Section 5.5.5*, this phase can easily be detected by the float voltage V_{float} , which is given in the data-sheet of the battery [Con11], and/or the charge controller.

5.5.5 Model Parameterization

To obtain the parameterization for the the charging and discharging model discussed in *Sections 5.5.3* and *5.5.4*, we perform a three-step procedure to extract the necessary information at design-time, i.e., in the laboratory, as explained in the following subsections. Note that the trace generation (*Section 5.5.5.1*), trace transformation (*Section 5.5.5.2*), and coefficient extraction (*Section 5.5.5.3*) must be performed only once for each battery type, i.e., chemistry, under consideration. The model then also applies to batteries of the same type, but different rated capacities. As we will show in *Section 5.7*, other battery conditions, e.g., age and capacity, or operating conditions, e.g., temperature, can be ignored during the offline phase. In *Section 5.6.1* we then elaborate on how the information collected during the offline phase enables runtime State-of-Charge approximation.

5.5.5.1 Trace Generation

To characterize the discharging and charging behavior of the battery, we first collect discharge and charge traces with different drain and charge currents respectively. Such a trace consists of fixed interval measurements of the system input voltage and constant drain current until the charge controller performs a low-voltage disconnect of the load (see *Section 5.5.1*). Optimally, a discharge (charge) trace is obtained for every possible drain (charge) current that the battery may be exposed to during deployment. However, this may not be feasible, or necessary. As will become clear in *Section 5.5.5.3*, at least three discharge traces should be obtained, where the magnitude of the drain currents $I_{drain,i}$ is chosen such that the traces represent the expected operational range of the system, i.e., $I_{sys,max} \leq I_{drain,i} \leq I_{sys,min}$. Similarly, the charge currents $I_{charge,i}$ should be chosen such that they cover the charge currents expected during deployment.

An illustration of manufacturer provided discharge traces for a Valve Regulated-Absorbent Glass Mat sealed lead-acid battery [Con11] at various discharge rates is shown in *Figure 5.2*. However, since the data provided by the manufacturer may be given for rates much higher than




Fig. 5.7: (top) Discharge traces of *BAT1* (see Table 5.3) and their approximations. (bottom) The resulting residuals of the approximation. Note that the voltages are relative to the cut-off voltage, i.e., $V_{cutoff} = 11.5V$.

required, the battery may have to be profiled with discharge rates appropriate for the particular system used. Here we assume the system introduced in Chapter 6, and fully described in [BYL⁺11].

5.5.5.2 Trace Transformation

After having obtained appropriate charge and discharge traces, we transform them from voltage as a function of time to Depth-of-Discharge as a function of voltage, which is defined for the interval $[0,1]$. This transformation is fully described in Sections 5.5.3 and 5.5.4.

5.5.5.3 Coefficient Extraction

In the following we discuss the approximation and interpolation coefficients for the charging and discharging models presented in Sections 5.5.3 and 5.5.4. For the parameterization we profiled *BAT1* (see Table 5.3) with relative loads of 0.00471, 0.01194, and 0.01638.

Discharge Model Coefficients

Figure 5.7 shows three traces of the discharge process of battery *BAT1* (see Table 5.3) and their quadratic approximations on the top graph (i.e., $n = 2$

in Equation (5.2)). The bottom graph in the figure shows the respective approximation error over the entire trace. From the bottom graph it is evident that the approximation achieves an accuracy of $\pm 2\%$ over the entire period, except for when the State-of-Charge is above 80%. This can be explained by the high dynamics during the transition from charging to discharging. While a cubic fit may improve the accuracy for a State-of-Charge larger than 80%, we found that it does not improve the model's performance appreciably. Moreover, a slightly higher error at a high State-of-Charge can be considered acceptable.

With the three traces, we can perform a quadratic approximation for the load interpolation, i.e., $m = 2$ in Equation (5.3). The respective coefficients are given in Table 5.1a. Note that if the quadratic interpolation accuracy is not sufficient, more than three traces must be generated, i.e., the number of traces $i = m + 1$. In the case under consideration, however, a binomial approximation is found to be sufficiently accurate.

Bulk Charge Model Coefficients

Figure 5.8 shows two traces with the same battery as before, i.e., BAT1, and two relative charging currents, i.e., $RC_1 = 0.03272$, and $RC_2 = 0.07246$. Since the traces are collected with known relative charging current, the corresponding slopes of the voltage profiles are used to extract the coefficients c_0 and c_1 of the first order approximation of \widehat{RC}_{bulk} , shown in Equation (5.4). The coefficients obtained with these traces are listed in Table 5.1b.

Absorption Charge Model Coefficients

The coefficients for the approximation of the decay constant $\widehat{\lambda}$ from Equation (5.7) are listed in Table 5.1c. As explained in Section 5.5.4.2, these coefficients are obtained by first finding the true decay constants $\lambda_i(RC_i)$ for each of the i charging traces with an exponential fit of the corresponding trace. Then we can define $\widehat{\lambda}(\cdot)$ with a linear approximation over the i true decay constants to obtain the coefficients d_1 and d_2 for Equation (5.7).

Float Charge Model Coefficients

As discussed in Section 5.5.4.3, the State-of-Charge for the float phase does not need to be approximated. However, as the float voltage V_{float} is dependent on the temperature T , we approximate it based on information from the battery's datasheet as shown in Equation (5.9). The temperature dependence implies that if the State-of-Charge during charging must be




Fig. 5.8: Linear approximation of the voltage slope of two different relative charging currents, i.e., $RC_1 = 0.03272$, $RC_2 = 0.07246$.

available, the system requires a temperature sensor that can measure the battery temperature.

$$V_{float}(T) = \max\{13000, 0.239 \cdot T^2 - 35.94 \cdot T + 14040\} \quad (5.9)$$

5.6 Online State-of-Charge Approximation

In this section we discuss how the models from the previous sections enable online approximation of the battery State-of-Charge. We first present a Finite-State Machine in *Section 5.6.1*, which is used to track the different charging and discharging phases, and compute the State-of-Charge accordingly. Then, in *Section 5.6.2* we briefly discuss measurement specific considerations.

5.6.1 Battery State Tracking

For tracking the behavior of the battery over time and updating the State-of-Charge according to the respective phase, we define a Finite-State Machine, of which the states and transitions are illustrated in *Figure 5.9*. The individual states correspond to the four charge and discharge phases explained in *Sections 5.5.3* and *5.5.4*. An additional UNKNOWN state is introduced to capture initial conditions. The individual states and triggers for each of the state transitions are listed in *Table 5.2*, and briefly described in the following.

Tab. 5.1: Coefficients for *Equations* (5.3), (5.4), and (5.7) found by interpolation of traces.

(a) Coefficients for Equation (5.3)					
Coefficient	Value	Coefficient	Value	Coefficient	Value
$b_{2,0}$	-8.321e-9	$b_{1,0}$	-0.07018	$b_{0,0}$	100
$b_{2,1}$	-0.002257	$b_{1,1}$	4.492	$b_{0,1}$	0
$b_{2,2}$	0.08133	$b_{1,2}$	-204.7	$b_{0,2}$	0

(b) Coefficients for Equation (5.4)		(c) Coefficients for Equation (5.7)	
Coefficient	Value	Coefficient	Value
c_0	0.005573	d_0	-0.003891
c_1	1.296	d_1	-0.1091

UNKNOWN. This is the starting state. While in this state, the system gathers information, i.e., system input voltage V_{sys} and drain current I_{sys} to determine the next state. The system monitors the evolution of the battery voltage and drain current for detecting the triggers for the state transitions shown in *Table 5.2*. If a dropping voltage tendency is observed, the system enters the DISCHARGE state, irrespective of the current. On the other hand, if an increasing voltage combined with a constant or increasing current is observed, the state machine transitions to the BULK state. Note that if the current decreases, the voltage will rise accordingly, hence a decision is deferred as it is not clear if the battery is being charged or discharged. The ABSORPTION state is entered if the voltage reaches the absorption voltage level V_{abs} , and the FLOAT state is entered if the measured voltage is less than the absorption voltage, but greater than or equal to the float voltage level V_{float} . Note that no State-of-Charge computations are performed in this state.

DISCHARGE. This state represents the discharging process of the battery. The state machine remains in this state as long as the battery is being discharged, i.e., while the voltage drops. In this state the system monitors the battery voltage and the drain current for detecting the start of the bulk phase, which is characterized by a monotonically increasing voltage without a reduction in drain current. Upon detection of the bulk phase, the system transitions into the BULK state. Note that in this discussion,

state machine transitions to the DISCHARGE state. The State-of-Charge in the BULK state is computed with *Equation (5.5)*.

ABSORPTION. The typical characteristic of the absorption phase is that the battery voltage stays relatively stable and does not keep increasing, as is the case for the bulk phase. The State-of-Charge in this state is computed with *Equation (5.8)*. Once the observed battery voltage drops from the absorption voltage threshold V_{abs} to the float voltage level V_{float} the state machine transitions to the FLOAT state. If the battery voltage falls below the float threshold voltage V_{float} , the battery is not being charged anymore, hence state machine transitions to the DISCHARGE state.

FLOAT. This is the final state of the charging process. In this state, the battery is fully charged and kept at 100% with a small charging current. If the observed voltage is below the float voltage level V_{float} , the battery is discharging and thus the system transitions into the DISCHARGE state. By definition, the State-of-Charge in this state is 100%.

5.6.2 Voltage Measurements

Since we assume an off-the-shelf setup depicted in *Figure 5.1*, the input voltage V_{sys} measured by the system does not correspond to the actual battery voltage. This is illustrated in *Figure 5.10*, which shows the electrical circuit equivalent of the assumed system architecture. The charge controller and power cable are assumed to be resistive loads. Note, that for profiling purposes we can also assume the system to be a constant resistive load. Now, since we need the battery terminal voltage to approximate the State-of-Charge in a setup independent manner, we must account for the voltage drops caused by the individual components, i.e., the charge controller and the power cable, as shown in *Equation (5.10)*. The unknown battery terminal voltage is denoted V_{bat} , while V_{sys} represents the measured input voltage, V_{cc} stands for the voltage drop due to the charge controller, and V_{cable} refers to the voltage drop caused by the power cable.

$$V_{bat}(t) = V_{sys}(t) + V_{cc} + V_{cable}(t) \quad (5.10)$$

Note that we assume V_{cc} to be constant, obtained from the charge controller's data-sheet or experimentally approximated. V_{cable} can be computed using the cable's specific resistance ρ , its length l and cross-area A , multiplied by the measured drain current I_{sys} , as shown in *Equation (5.11)*.

$$V_{cable}(t) = \rho \cdot \frac{l}{A} \cdot I_{sys}(t) \quad (5.11)$$




Fig. 5.10: Simplified model of the voltages and currents of the system platform.

5.7 Evaluation

Next we present performance result of the proposed State-of-Charge approximation approach when used with an off-the-shelf power subsystem and high energy consumer, e.g., wireless sensor network basestation [MOH04], or applications with high-power sensors and increased duty-cycles [BSB12].

5.7.1 Experimental Set-up

5.7.1.1 Test Set-up

For evaluating the proposed approach, we carry out discharge tests in our lab with three Valve-Regulated Absorbent Glass-Matt (VR-AGM) batteries [Con11] listed in *Table 5.3*. A custom platform, of which details are given in [BYL⁺11], presents the load and executes the logic described in *Section 5.6* with the parameters shown in *Table 5.4*. To verify that batteries of the same capacity yield the same results, a discharge test is performed on two identical batteries. The differences are found to be insignificant, hence only one set of results is presented in the following.

5.7.1.2 Model Parameters

General equations for obtaining the Dept-of-Discharge as a function of the battery voltage and relative load have been developed in *Section 5.5*. The model parameterization has been discussed in *Section 5.5.5*. *Table 5.1* list the respective coefficients, which have been found by analyzing discharge traces of a single battery, i.e., *BAT1* (see *Table 5.3*), which was profiled at 22 °C as described in *Section 5.5.5*.

Tab. 5.3: Type, capacity, and condition of the batteries [Con11] used.

Name	Type	Capacity (C/20)	Condition
BAT0	GPL-1400T	43 Ah	Old
BAT1	GPL-U1T	33 Ah	New
BAT2	GPL-1400T	43 Ah	New

5.7.1.3 Evaluation Baseline

In order to evaluate the method's ability to adapt to conditions that are not explicitly modeled, such as varying size, condition, etc., a number of discharge tests are carried out and discussed in the next section. Discharge tests under controlled conditions are considered to yield accurate State-of-Charge approximation [PPJ01, Con11]. Before starting each test, the respective battery is fully charged and allowed to rest for 4 hours.

To obtain a baseline for performance comparison and validation of the assumption about available battery capacity made in *Section 5.5.1*, an external coulomb counter [Max98] is placed between the charge controller and the battery. The coulomb counter measures the current $i(t)$ drawn from the battery during the discharge tests discussed in the next section. These measurements are converted to the baseline State-of-Charge $SoC_{CC}(t)$ by assuming that the battery is charged to full capacity C at the start of the discharge test, i.e., $SoC_{CC}(0) = C/C$, and then defining the baseline State-of-Charge as shown in *Equation (5.12)*, where Δt refers to the sampling interval, and t_{cutoff} is the time when the charge controller performs the low-voltage disconnect.

$$SoC_{CC}(t) = \frac{C - i(t) \cdot \Delta t}{C} \quad \forall t[0, t_{cutoff}] \quad (5.12)$$

To minimize the errors introduced by the coulomb counter, it is calibrated according to its datasheet.

5.7.2 Experimental Results

In the discussion so far, we have ignored battery inherent inefficiencies. However, when operated under varying, and uncontrollable conditions, batteries exhibit non-linearities that cannot be neglected. Therefore, we first evaluate the performance of the State-of-Charge approximation with a battery exposed to varying temperatures in *Section 5.7.2.1*. As a second

Tab. 5.4: Sampling rate and filter specification for smoothing of physical measurements.


Description	Value
Sampling rate	30 sec.
Filter type	EWMA
Filter length	10

step, in *Section 5.7.2.2* we investigate the model’s ability to adapt to aging batteries. Then, in *Section 5.7.2.3* we present results obtained with varying loads. Finally, in *Section 5.7.2.4* we discuss the model’s accuracy in predicting remaining runtime. Note that the goal of the experiments is to evaluate the performance of the proposed approach, and *not* the battery behavior under the different conditions.


5.7.2.1 Temperature Effects

Temperature variations have a strong impact on the battery’s State-of-Charge as they affect the apparent capacity [RVR03, Con11]. Low temperatures cause a reduction of the electrochemical activity, which leads to a (temporary) reduction of charge stored in the battery (relative to its rating), and to some extent the opposite is true for high temperatures. Since the battery parameters were obtained at 22 °C, with the first experiment we evaluate the model’s ability to adapt to lower temperatures. To this end, *BAT1* is cooled down and kept at −6 °C while being discharged with a constant relative load of 0.0103. This specific temperature is selected as it represents a typical average temperature during winter months at a particular deployment location of interest, see *Section 6*. *Figure 5.11* and *Table 5.5* show the results, which we briefly discuss in the following.

The upper graph in *Figure 5.11(a)* shows the State-of-Charge approximation by the proposed approach and the ground-truth at constant −6 °C, and for reference at 26 °C. The lower graph shows the State-of-Charge approximation error, which is defined as $SoC_{model}(t) - SoC_{CC}(t)$, i.e., the difference between the State-of-Charge indicated by the model and that inferred by the baseline. At −6 °C the maximum overestimate of 3.54% occurs only in the very beginning but quickly drops during the first 6 hours until reaching a maximum underestimate of 1.17% after 12 hours. After 25 hours, and for the remainder of the test, the error varies between roughly 0.3% and 2% with a mean deviation from the ground-truth of 0.85%.



(a) State-of-Charge approximation (top) and error (bottom).



(b) Runtime prediction (top) and prediction error (bottom).

Fig. 5.11: Constant power discharge test at constant 26°C and -6°C with battery *BAT1*.

The reason for the initial overestimate is likely due to the increased rate at which the voltage drop occurs due to the low temperature. This causes the model to overestimate the State-of-Charge until the battery has adjusted to the load, at which point the rate of change in the voltage profile slows down and reflects the actual condition. The consistent overestimate in both cases is likely because of considering too little voltage drop across the charge controller and wires (see *Section 5.6.2*).


Another inefficiency that is highly temperature dependent and common to all batteries is referred to as self-discharge. This phenomenon causes the battery to discharge at a type and temperature specific rate [AUDT10] even when no load is connected. We disregard this effect as it is assumed that an off-the-shelf charge controller presents a load that is always larger than the battery's self-discharge.

In summary, the temperature, and – as we will show in the next section – aging effects are indirectly considered through changes in terminal voltage that occur due to temperature variations without a corresponding change in drain current. The change in voltage induced by these effects causes the model to yield a correspondingly lower or higher State-of-Charge. It is important to note that 0% State-of-Charge is not affected by temperature because the charge controller disconnects the load at a temperature independent cut-off voltage.


5.7.2.2 Aging Effects

Each charge/discharge cycle causes irreversible chemical reactions within the battery. This causes a gradual decrease in the maximum charge that the battery can store and deliver [BKN02], and is referred to as capacity loss. Conventionally, a battery has reached its end-of-life when it fails to store and deliver a certain percentage (usually 80%) of its nominal capacity [Con11]. Since capacity and voltage are directly related [PPJ01], an old battery will exhibit a lower terminal voltage (compared to a new battery) due to lower electrochemical activity. Therefore, identically to temperature effects discussed in the previous section, the model indirectly considers aging effects via its impact on terminal voltage.

Figure 5.12 and *Table 5.5* show the results of the discharge test with *BAT0*, which has reached its end-of-life. For reference, the same discharge test is carried out with *BAT2*, which is of the same type and capacity, but in new condition. The initial capacity is determined using a coulomb counter, which indicates that only 52% of the nominal capacity (equivalent to roughly 65% of apparent capacity) can be withdrawn when *BAT0* is fully charged. As is evident from the top most graph in *Figure 5.12(a)*, this agrees very well with the initial State-of-Charge approximated by the model. The maximum and mean deviation from the



(a) State-of-Charge approximation (top) and error (bottom).



(b) Runtime prediction (top) and prediction error (bottom).

Fig. 5.12: Discharge test with aged (*BAT0*) and new (*BAT2*) batteries.

ground-truth of -4.49%, and -1.69% respectively may seem rather high. However, the negative values represent an underestimate, which is maintained almost over the entire range. This may be viewed as a conservative indication of State-of-Charge, and, depending on application, may even be desirable. It is worth noting that, except for in the beginning, both discharge tests yield very similar profiles, which validates the method's applicability to varying battery conditions. Unfortunately, the model parameters, which are obtained with a battery size equivalent to *BAT1*, tend to introduce some error when used with the larger batteries. The effect is visible for both cases by the dip in the profile shortly after starting the test.


5.7.2.3 Load Variations

In practical applications, the system typically does not exhibit constant power dissipation, hence this section presents the model's performance when exposed to varying drain currents. Load changes frequently occur in real systems, due to e.g., duty-cycling. It is well known that a battery behaves non-linearly when discharged at different rates [Peu97, DAS06]. These rate dependent effects are known as rate-discharge and recovery effects.


Figures 5.13 and 5.14 and Table 5.5 present the results of a discharge test with bimodal load changes. As illustrated in the lower graph in Figure 5.13, which shows the voltage profile, the *high-power* mode presents a maximum relative load of 0.0170 and is active 35% of the time, while the *low-power* mode presents a relative load of 0.0055 for the remaining 65%. In this discharge test, the rate-discharge and recovery effects become visible. With variations in the load, the battery is allowed to recover some of the charge regularly. Since this affects the battery's terminal voltage, the model tends to continuously fluctuate between slightly under- and overestimating the State-of-Charge. Towards the end of the discharge test, the maximum overestimate reaches 4.98%. Nevertheless, despite the load changes and their visible effects on the voltage profile, the State-of-Charge exhibits a relatively linear tendency with a mean deviation from the expected value of 3.02%.

5.7.2.4 Lifetime Estimation

The ability to predict the remaining runtime is intuitively more beneficial to the system than mere indication of battery State-of-Charge. For this reason, in this section we investigate if the approximated State-of-Charge can be easily used to predict remaining runtime. In addition to the instantaneous State-of-Charge and drain current, the manufacturer



(a) State-of-Charge approximation for bimodal load.



(b) Voltage profile for bimodal load

Fig. 5.13: State-of-Charge approximation (top), and voltage profile (bottom) for bimodal load.

specified nominal capacity C in mAh is used as basis for computing the remaining lifetime L , as shown in *Equation (5.13)*.

$$L(t) = SoC(t) \cdot \frac{C}{I_{sys}(t)} \cdot \alpha \quad (5.13)$$

The scaling factor α is introduced to account for rate-discharge and recovery effects with variable loads. Therefore, all but the experiment with varying loads use $\alpha = 1$. In the more interesting case with varying loads, rate-discharge and recovery effects have an impact on the assumed capacity, which requires an appropriate scaling to yield acceptable life-time prediction. For the results shown in *Figure 5.14*, $\alpha = 0.7$ has been experimentally found to yield the lowest average lifetime prediction error.

It is to be expected, however, that α does not behave linearly with the discharge duty-cycle. To verify this, we carried out another set of tests




Fig. 5.14: Runtime prediction (top) and prediction error (bottom) for bimodal load.

with a duty-cycle of 50% and periods of one ($\alpha = 0.85$), two ($\alpha = 0.88$), and four hours ($\alpha = 0.88$). This illustrates the fact that the rate at which the load varies, affects the actually available charge in the battery [BKN02]. The implication is that for a simple lifetime prediction approach like *Equation (5.13)*, α may have to be learned online if the duty-cycle is not known in advance.

Another complication arises due to the discrete, and potentially noisy current measurements $I_{sys}(t)$ to compute the lifetime. This is illustrated in the upper graph in *Figure 5.14*, which shows the runtime predicted by the model (see *Table 5.4* for smoothing filter parameters). The smooth curve represents the predicted lifetime with a running average for $I_{sys}(t)$, while the dashed curve represents the prediction when using the instantaneous measured current instead. For reference, the lifetime prediction with absolute average, i.e., averaging over all measurements observed so far, for $I_{sys}(t)$ is represented with the dash-dotted line. As is to be expected, when using instantaneous measurements, the prediction follows the voltage profile closely (which follows the drain current) and so introduces relatively large errors. On the other hand, the averaging

operation smoothens the measurements, which allows for predicting the expected lifetime much more closely. Depending on application, either of the approaches may be favorable.

Table 5.5 and the lower graph in *Figure 5.14* show the prediction error for the running, and absolute average respectively. It would not be fair to compare the instantaneous lifetime prediction to the actual runtime, which is a product of the discharge dynamics. Nevertheless, despite initially exhibiting large deviations from the actual runtime by up to -12.29% , the average error over the entire test is found to be only -0.85% . This is a bit skewed due to the consistent, and quite significant underestimation of the lifetime during the first half of the test, which exhibits high dynamics in the load variation. However, for the second half, the mean error is roughly 1.9% with an absolute maximum and minimum of 4.52% and 0.04% respectively.

In summary, especially the rate-discharge and to some extent the recovery effects have a large impact on lifetime prediction when load changes occur. As expected, these effects are less pronounced when a constant load is present. Overall, the simple lifetime prediction approach from *Equation (5.13)* is able to quite accurately approximate the lifetime with a maximum underestimation of 12.29% for the variable load test, and a maximum overestimation of 9.21% for the test with the aged battery.

5.7.2.5 Charge Approximation

Approximation of the State-of-Charge during the charging process has been tested under optimal conditions and works very well. However, when dealing with real-world conditions where the solar panel's performance varies due to e.g., clouds, the state machine introduced in *Section 5.6.1* has difficulties properly tracking the states, which in turn causes the approximation accuracy to suffer. Considering the good performance during discharge – which is the time when State-of-Charge approximations are most useful – this is not considered an issue. Moreover, in the case of solar energy harvesting, the system tends to spend the majority of time in the discharging state.

Tab. 5.5: Results of the State-of-Charge approximation performance shown in *Figures 5.11 through 5.14*.

(a) State-of-Charge error in % ($SoC_{model} - SoC_{CC}$).						
Test	Figure	Maximum	Minimum	Mean	Variance	Std.
26 °C	5.11(a)	2.98	0.33	2.0	0.44	0.66
-6 °C	5.11(a)	3.54	-1.1	0.85	0.96	0.98
Aged	5.12(a)	2.41	-4.49	-1.69	3.24	1.8
New	5.12(a)	2.04	-5.32	-1.81	5.96	2.44
Var. RL	5.13	4.98	-0.79	3.02	4.53	2.13
Mean	–	3.19	-2.27	0.47	3.03	1.6

(b) Lifetime error in % of actual runtime.							
Test	Figure	Maximum	Minimum	Mean	Var.	Std.	α
26 °C	5.11(b)	3.85	-1.38	1.4	2.95	1.72	1.0
-6 °C	5.11(b)	7.54	-1.34	1.95	2.87	1.7	1.0
Aged	5.12(b)	9.21	-3.52	-0.07	9.52	3.09	1.0
New	5.12(b)	7.04	-1.13	0.96	4.63	2.15	1.0
Var. RL	5.14	3.8	-12.29	-0.85	18.83	4.34	0.7
Mean	–	6.29	-3.93	0.68	7.76	2.6	–

5.8 Measurement Considerations

In this chapter we devised a light-weight HW/SW approach to battery State-of-Charge approximation that does not depend on special-purpose hardware. Nevertheless, since the approach relies on voltage measurements, at least a voltage sensor is required that can measure the unregulated system input voltage with sufficient resolution and accuracy. The drain current can be obtained by profiling, as is done in [CPBA12, DOTH07, KPA⁺10]. However, the instantaneous State-of-Charge approximation can be improved by employing a current sensor to

measure the true system drain current. If the State-of-Charge should also be tracked during charging of the battery, a temperature sensor to measure the battery's temperature must also be available because the float voltage is temperature sensitive, see *Section 5.5.5.3*. These sensors are generally present on wireless sensor network motes for health monitoring purposes, and are therefore not considered special-purpose hardware.

During discharging of the battery, the proposed algorithm is stateless and only considers the present measurements to compute the State-of-Charge. This means that the model's accuracy is independent of the sampling period, and the measurements for health monitoring of the respective application can be conveniently re-used whenever they are scheduled anyway. This eliminates measurement overhead and reduces possible interference with normal system operation. However noise on the sensor readings may affect the solution quality. For this reason, at runtime, an exponentially weighted moving average (EWMA) filter with length 10 is used to smoothen all physical measurements (see *Table 5.4*).

For the charging process, the State-of-Charge approximation depends both on the present measurement and previous State-of-Charge. In this case it is clearly beneficial to use a sampling period appropriate to capture the dynamics of the energy source used. With solar harvesting the source tends to fluctuate slowly, so a sampling period of a few minutes can be considered sufficient.

With an off-the-shelf set-up assumed, as illustrated in *Figure 5.1*, the system input voltage does not correspond to the actual battery voltage. However, to approximate the State-of-Charge in a set-up independent manner, the battery's terminal voltage is required. Therefore, voltage drops caused by the charge controller and the power cable, or any other consumers must be explicitly accounted for, as described in *Section 5.6.2*. Finally, it is worth noting that for this algorithm no State-of-Charge updates can be made during system sleep periods, unless supervisory circuit with the ability to read the sensors and execute the Finite-State Machine introduced in *Section 5.6.1* is available.

5.9 Chapter Summary

In this chapter we presented a light-weight approach to battery State-of-Charge approximation that relies fully on closed-loop voltage and optional drain current measurements and does not require any special purpose hardware. Instead it leverages the behavior of the battery under load, which is extracted from discharge traces. Therefore, it incurs only limited implementation effort to adapt for a variety of systems. The proposed approach achieves state-of-charge approximation with up to 95% accuracy

even under varying operating conditions. Minimal operational overhead make this approach suitable for a wide variety of low-performance embedded computing systems that leverage off-the-shelf energy harvesting set-ups.

6


Case Study: X-Sense GPS System

6.1 Introduction

In this chapter we demonstrate the feasibility of the theoretical models formulated in *Chapters 2-5* for realization with a realistic system. We use the context of the *X-Sense* project [BBF⁺11] as a case study, and quantify the benefits of the proposed dynamic power management schemes from *Chapters 3 and 4* in terms of long-term operation and increased system utility for the target application.

Motivation


The *X-Sense* project's primary goals are to (i) apply WSN technology to enable geoscientific characterization and quantification of cryosphere phenomena, e.g., [WBB⁺13, WGGP15], and their transient responses to climate change, and (ii) investigate the feasibility of WSN technology for an early-warning system against destructive events triggered by these phenomena. An integral requirement of the geoscientific aspects of the project is the ability to accurately track variable slope movements that range from a few centimeters per year to a few centimeters per day. Moreover, considering that the intended deployment site is difficult to access, we require a system that can reliably and autonomously provide positioning with sub-centimeter accuracy over time periods on the order of multiple years.



(a)



(b)



(c)

Fig. 6.1: (a) System block diagram partitioned into harvesting subsystem (solar panel), storage subsystem (battery and charge controller) and the load to be supported (wireless sensing system). The flow of generated energy is indicated with dotted lines, while the dashed lines represent the energy consumed. (b) Picture of the *X-Sense CoreStation* evaluation platform installed at the high-alpine deployment site. (c) Picture of the *CoreStation* evaluation platform, indicating the various system components.

To investigate the feasibility of the above two goals, a number of in-situ evaluation and experimentation platforms, called *CoreStation* (see *Figure 6.1c*), enhanced with commercially available single-frequency GPS receivers (μ -blox LEA-6T) are installed at the high-alpine deployment site (see *Figure 6.1b*). The raw L1 GPS data collected by these systems is sent to our back-end servers, after which the data is differentially post-processed [LGB⁺11, BSB12, SGL⁺14] to yield the positioning accuracy required by the geoscientific project partners [WBB⁺13]. To achieve an acceptable solution accuracy that allows tracking of slope movements with sub-centimeter resolution, the GPS receiver must be active for time periods exceeding 2 hours (see *Section 6.2.2*). Therefore, given the long sampling intervals, coupled with the GPS receivers' high power demands [JCDS10], and the need for unattended operation over multiple years, energy harvesting technology is indispensable to power the *CoreStation* platform.

In summary, the project requirements and characteristics of the experimentation platform present an ideal scenario for the validation of the assumptions and models in the preceding chapters, and the evaluation of the dynamic power management schemes from *Chapters 3* and *4* under real-world conditions.

6.2 System Model

6.2.1 System Architecture

The system architecture is illustrated in *Figure 6.1a*. From a high-level perspective, the system consists of the solar energy harvesting and energy storage subsystems, collectively referred to as power subsystem, and the load, i.e., wireless sensing platform, to be supported.

To satisfy the project requirements from *Section 6.1*, we leverage a custom-built, feature-rich hardware platform together with an extensible middle-ware, briefly described in *Section 6.2.1.1*. With this platform we trade off system flexibility, observability, and accessibility for relatively high power dissipation (i.e., roughly 6 Watt at full performance). Since batteries alone could not support the high power requirements over extended periods, the system is powered by the solar energy harvesting system discussed in *Section 6.2.1.2*.

6.2.1.1 Wireless Sensing Platform: *CoreStation*

A picture of the deployed X-Sense *CoreStation* is shown in *Figure 6.1b*, and *Figure 6.1c* shows a close-up of the *CoreStation*'s hardware platform enclosed in a water-proof die-cast box with EMP-protected antenna, power,

and debug connectors. The following two sections provide an overview of the hard- and software architecture.

Hardware

To accommodate for the evaluation of diverse application scenarios, the hardware architecture is partitioned into three layers: power and interfacing, processing, and networking.

Power and Interfacing. The Baseboard accepts 5.5V-18V unregulated input and generates multiple software controlled and monitored power domains. In addition to onboard temperature, humidity, voltage and current sensors, two USB interfaces, one configurable serial interface, and one I^2C bus for connecting external sensors are available. The Baseboard further provides interfaces and communication planes for the Gumstix embedded PC [gum], a TinyNode [DFFMM06] (or compatible) mote, and a GSM modem. Accurate time across power cycles is maintained by a battery-backed real-time clock. Finally, for on-site maintenance, a serial console can be accessed without the need to open the case.

Processing. The Gumstix verdex pro XL6P provides the main processing element of the *CoreStation*. The 600MHz XScale CPU, together with 128MB of RAM, and 8GB of Flash memory provides enough resources to run a Linux operating system, with most of the features and utilities available on desktop distributions. In the current set-up we further utilize a TinyNode184 mote, which features an MSP430 low-power processor, running the TinyOS operating system [LMP⁺05]. The Gumstix and TinyNode communicate via the TinyOS serial protocol over UART.

Networking. Our setup employs two redundant high-bandwidth links, i.e., WLAN and GSM, which provide reliable access and visibility into the system. The *CoreStation* further supports ethernet or satellite links. The TinyNode's radio permits integration into ultra low-power Dozer [BVW07] networks. Using special Dozer beacons, commands to either of the processing elements can be sent from remote via the ultra low-power network.

Software

The custom software framework, called *BackLog* runs on top of the Linux OS. It provides a plug-in API, power control module, scheduler for execution of individual plug-ins, and optionally handles connection to the GSN [AHS07] backend servers. Built-in monitoring and statistics, together with high-level language support and well-established utilities (ssh, snmp, ntp, etc.) prove to be invaluable in rapidly implementing

functionality and analyzing system performance trade-offs. *BackLog* further provides communication interfaces with, and control over hardware components. The ability to interact with individual components, such as monitoring or reprogramming the TinyNode, polling sensors, or enable/disable wireless links and other system components, extends the system functionality to an on-site testbed of low-power, TinyNode based sensor nodes.

System Limitations

Although proven to be a very powerful and flexible tool for in-situ experimentation, the evaluation platform has one major limitation: it can not measure the current generated by the panel I_{pv} , and only monitor the system input voltage V_{sys} and current drawn I_{sys} (see *Figure 6.1a*). However, as discussed in *Section 3.3*, an approximation of the generated energy $E_{real}(d)$, which can be derived from I_{pv} , is necessary for the dynamic power management algorithms from *Chapters 3* and *4* to function. In *Section 6.3* we explain how the lack of appropriate hardware support can be circumvented.

6.2.1.2 Power Subsystem

As illustrated in *Figure 6.1a*, and mentioned in *Section 3.2.1*, we assume a harvest-store-use architecture [SK11], which is enforced by the employed charge controller. This means that the battery is not bypassed when the panel generates surplus energy, i.e., when the battery is full. The implication is that inefficiencies due to battery charging and discharging are incurred even if the panel could directly power the load.

We consider two installations that are identical with respect to *technology* and *system* parameters (see *Figure 2.1* in *Chapter 2*): a 30 Watt mono-crystalline solar panel (cleversolar CS-30) with a solar cell area of $0.1725m^2$, a SunSaver SS-6L PWM charge controller, a Lifeline AGM battery with a nominal capacity of 54Ah, and the *CoreStation* wireless sensing system described above with a power dissipation of approximately 6 Watt at full performance. *Table 6.1* lists the relevant *deployment* parameters that differ between the two installations.

Note that, according to our model in *Chapter 2*, neither of the above two configurations is ideal with respect to the dimensioning of the energy harvesting power subsystem. The panel and battery are poorly matched: the battery is under-provisioned for the given load (the deployed battery is only dimensioned for 10 days of operation at a duty-cycle of 30%, i.e., average power dissipation of roughly 2 Watt, and hence not large enough to bridge extended periods of snow cover during winter), and the panel is

Tab. 6.1: Name, time-period, coordinates, and solar panel orientation (ϕ_p) and inclination (θ_p) angles of the deployed systems used in the case study.

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

over-provisioned for the battery employed (resulting in available energy being wasted during summer time). Nevertheless, using the capacity planning algorithm from *Section 2.3.1* and assuming environmental conditions specified by $\Omega = 0.6$, we find that this harvesting configuration should be able to support 7.5 hours (i.e., $DC = 31.25\%$) of daily system operation for the DH location, and 6.7 hours for the GG location (see *Table 6.1*), as long as any one period without harvesting opportunities does not exceed 10 days.

In order to reduce the likelihood of a low-voltage disconnect due to the under-provisioned battery, we adjust the parameterization for the dynamic power management scheme from *Chapter 3* to use a history window size $W = 10$ days (see *Sections 3.3.1* and *3.6.3*) to match the period that can be bridged by the battery. For the approach from *Chapter 4* we use the parameterization given in *Section 4.6*.

6.2.2 System Utility

In the previous chapters we assumed the system utility to be directly proportional to the duty-cycle, which the system can achieve. Here we refine it appropriately for the aforementioned case study.

As introduced in *Section 6.1*, the application scenario [BBF⁺11, WBB⁺13] relies on a differential GPS processing algorithm, whose acceptable error performance requires periodic sampling of the GPS receiver (μ -blox LEA-6T) over at least two consecutive hours per day. This is illustrated in the top graph in *Figure 6.2*, which shows the processing algorithm's error performance as a function of the measurement duty-cycle [Lim]. From the figure, it is evident that the minimum acceptable error, i.e., $e(DC) \leq 8\text{mm}$, requires a duty-cycle of at least 8.33%, which corresponds to a minimum of two hours of continuous sampling per day [BSB12], above which the error decreases exponentially. Note that the error is undefined for duty-cycles lower than 8.33%, as this is below the minimum requirement imposed by the processing algorithm [Lim].




Fig. 6.2: Error performance (based on [Lim]) of the processing algorithm versus duty-cycle (top), and derived system utility versus duty-cycle (bottom). Note that the error performance is undefined for a duty-cycle lower than 8%.

The application under consideration [WBB⁺13] directly benefits from increased temporal resolution, e.g., to characterize sub-daily process variations, hence, sampling over longer time intervals increases the system utility. Therefore, we define the system utility, as shown on the bottom graph in Figure 6.2, to be $U(DC) = 1 - e_{norm}(DC)$, where $e_{norm}(DC)$ represents the normalized error performance $e(DC)$ shifted by an offset to reach the maximum utility at a duty-cycle of 100%. This is a realistic definition of system utility for many application scenarios, e.g., [DEM⁺10, CB10].

6.3 Implementation Details

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

6.3.1 Circumventing Limited Measurement Support

For a number of design decisions, e.g., to reduce failures due to external wiring, the hardware platform is not designed to provide measurements of the energy generated by the solar panel, or the energy actually flowing into the battery. However, the dynamic power management schemes introduced in this work rely on an approximation of the harvested energy on a given day. In order to approximate the harvested energy without appropriate hardware support, we leverage the battery State-of-Charge algorithm proposed in Chapter 5. At an absolute minimum, this algorithm

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

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

Since we are interested in the energy actually generated by the panel, and not its theoretical daily maximum, we scale $\widehat{H}_{max}(d)$ by a factor ζ to approximate $E_{real}(d)$. The scaling factor ζ accounts for the fact that, with the given configuration, a certain fraction of the energy generated by the panel is wasted. We approximate this scaling factor as shown in Equation (6.1) by considering only the days in the past on which the State-of-Charge (SOC) approximation is below 80%.

$$\zeta_d = \text{mean} \left(\zeta_{d-1} + \frac{G(d)}{\widehat{H}_{max}(d)} \right), \text{ if } SOC(d) \leq 80\%, \forall d \quad (6.1)$$

The quantity $G(d)$ is the daily energy generation approximated using measurement of the system's current drain, and the State-of-Charge approximation, as given in Equation (6.2). Note, V_{bat} is the operating voltage, i.e., $V_{bat} = 12V$ in our case, and B is the apparent battery capacity, i.e., $B = \eta_{out} \cdot B_{nom}$.

$$G(d) = \frac{\widehat{E}_{out}(d)}{\eta_{in} \cdot \eta_{out}} + \frac{(SOC(d) - SOC(d-1)) \cdot B \cdot V_{bat}}{\eta_{in}} \quad (6.2)$$

6.3.2 LT-ENO Algorithm Modification

The *LT-ENO* dynamic power management algorithm discussed in Chapter 3 was designed for enabling uninterrupted long-term operation as the primary goal. The algorithm assumes an appropriately provisioned

battery and solar panel, as obtained with the capacity planning from *Chapter 2*. However, as already discussed, the configuration under consideration is suboptimal, therefore, in the following we discuss an adjustment that aims at improving the algorithm's performance with the given power subsystem configuration.

Since the panel is over-provisioned, it will likely be able to generate significantly more energy during periods of surplus than what can be stored in the battery, and used by the system (see *Chapter 2*). Hence, during periods of surplus (i.e., on the interval $[d_0, d_1)$, see *Figure 3.3*), we scale the duty-cycle computed by the DPM algorithm by a factor $\psi = 2$. This value has been determined experimentally, which suggest that the exact value for ψ depends on the ratio of energy consumption and generation. Furthermore, to reduce excessive usage of the energy stored in the battery during critical times, we introduce a guard time $T_g = 30$ days, which has been defined based on experiments. Therefore, the modification will only be used on day d , if and only if the following inequality holds: $d_0 + T_g < d < d_1 - T_g$. Recall that d_0 and d_1 are computed dynamically (see *Algorithm 1* in *Section 3.3.1*).

6.4 Performance Evaluation

In this section we present and discuss the results obtained from executing the dynamic power management algorithms from *Chapters 3* and *4* on the system introduced in *Section 6.2* so to expose it to realistic settings.

In order to verify the simulation framework's applicability for investigation of different algorithm parameterizations, we first briefly discuss our validation methodology: We recorded traces¹ of battery State-of-Charge, harvested energy, and load of the wireless sensor platform described in *Section 6.2.1*. We then set the parameters of our simulation framework to represent this system and compared the recorded State-of-Charge trace to the output of the simulation framework for the same period. We observed a Mean Absolute Percentage Error [AC92] between measured and simulated battery State-of-Charge of 5.45% over the entire trace. Therefore, we can conclude that a simulation with other algorithmic parameterizations will exhibit similar low deviations.

6.4.1 Evaluation Results

In the following we first evaluate the dynamic power management scheme from *Chapter 3*, with and without the modification discussed in

¹All datasets used in this work are available at <http://data.permasense.ch>


Section 6.3.2, to which we refer as $LT - ENO^+$ and $LT - ENO$ respectively. Then, in Section 6.4.1.2 we evaluate the performance of the optimal power management scheme from Chapter 4, to which we refer as OPT .

6.4.1.1 $LT - ENO$ and $LT - ENO^+$ Performance Results

The top graphs in Figures 6.3 and 6.4 show the daily energy $E_{real}(d)$ observed by the respective system, together with the modeled energy expectation $E_{in}(d)$ and the modeled energy consumption $E_{out}(d)$ for the two deployed systems (see Section 6.2.1) over almost 700 days. Recall that $E_{in}(d)$ is derived from the astronomical model at design-time, which, depending on its parameterization, can be an optimistic or pessimistic approximation. The bottom graphs show the static duty-cycle ($Static_{exp}$) that the system is expected to support based on design-time model assumptions, and the dynamically computed duty-cycle $LT - ENO^+$, which incorporates the modification discussed in Section 6.3. For reference, we also show the duty-cycle if this modification were to be disabled (referred to as $LT - ENO$), and the static duty-cycle $Static_{max}$ that may actually be supported by the respective energy harvesting configuration when no dynamic power management is employed. Note that for the latter, we assume that perfect knowledge of the true energy conditions are available at design time.

From the top graphs in Figures 6.3 and 6.4, we first note that the expected harvesting opportunities were significantly overestimated with the selection of the environmental parameter $\Omega = 0.6$ (see Section 6.2.1). A more appropriate, and safer parameterization would be $\Omega = 0.83$ for location DH, and $\Omega = 0.8$ for location GG. This results in a supported duty-cycle $Static_{max} = 12.9\%$ for DH, and $Static_{max} = 10.8\%$ for GG, assuming all other parameters are unchanged.

By inspection of the bottom graphs in Figures 6.3 and 6.4 we note that $LT - ENO$ stays between the expected duty-cycle $Static_{exp}$ based on model assumptions and the supported static duty-cycle $Static_{max}$ for the majority of the time. This clearly shows the proposed algorithm's ability to dynamically adapt the performance level in response to deviations from the modeled expectation, without risking battery depletion. As is evident from Table 6.2, the $LT - ENO$ approach eliminates power outages due to overestimating actual conditions, as experienced by the static approach, i.e., $Static_{exp}$, while simultaneously improving both system utility and energy efficiency. In fact, for the DH location, $LT - ENO$ achieves over 43% improvement in average system utility when compared to $Static_{max}$. The energy efficiency, i.e., the ratio of used and harvested energy, improves by 66% over $Static_{max}$. For the GG location, the improvements are even more substantial.

**Fig. 6.3:**

Uninterrupted operation over 700 days for the energy harvesting system at location DH.


(top) $E_{in}(d)$ refers to the offline model, which, in this case, consistently overestimated the true conditions. In general, $E_{in}(d)$ can be an optimistic or pessimistic approximation, however. $E_{out}(t)$ refers to the expected energy consumption in Wh , while $Static_{exp}$ is the corresponding duty-cycle.

(bottom) $Static_{max}$ is the truly supported duty-cycle, assuming the actual conditions $E_{real}(t)$ are known at design-time. $LT - ENO^+$ and $LT - ENO$ refer to the duty-cycle achieved by the dynamic power management scheme from Chapter 3 with and without the modification discussed in Section 6.3.2. $SOC(t)$ refers to the battery State-of-Charge, i.e., the battery's fill-level.

The results for the GG location, depicted in *Figure 6.4* and summarized in *Table 6.2* show that, although the system achieves zero downtime, the minimum achieved duty-cycle is below the minimum duty-cycle required by the end-user application (see *Section 6.2.2*). The resulting periods with zero utility are due to time periods with no harvesting opportunities that significantly exceed the time period the battery can bridge, i.e., 10 days. As discussed in *Section 3.3.2*, zero input conditions trigger an emergency mechanism, which overrides the duty-cycle computed by our algorithm due to lack of energy input. The effect of the emergency mechanism is clearly visible on the bottom graph in *Figure 6.4* on the intervals [31, 41], [253, 262], and [601, 619], all of which are longer than, or equal to ten days. Once the emergency mechanism is active, the resulting duty-cycle may be below the minimum required by the processing algorithm until a non-zero energy input is observed. Although experiencing days with zero energy input at the DH location in *Figure 6.3* as well, these instances are not as prominent in the dataset.

Next we consider the modified approach, referred to as $LT - ENO^+$, which differs from $LT - ENO$ only during periods of energy surplus, i.e., summer. The effect of the modification discussed in *Section 6.3.2* is clearly visible by the sudden step in duty-cycle around days 300, 400, and 670 for DH in *Figure 6.3*, and days 290, 430, and 650 for GG shown in *Figure 6.4*. From the results in *Table 6.2* we note that this modification does not affect the minimum utility, but improves the average utility and the system's energy efficiency. In fact, for the DH location, $LT - ENO^+$ achieves an improvement in average utility of 15%, and an impressive 80% in energy efficiency, when compared to $LT - ENO$. For the GG location, the improvement amounts to 14% and 61% respectively.

The performance improvements due to the efficient use of the power subsystem are ultimately expected to translate into increased system utility for the end-user. *Figure 6.5* shows the histogram of system utility (defined in *Section 6.2.2*) achieved by the four different approaches. The static approach with $Static_{max}$ achieves a constant, but clearly the lowest utility. As is to be expected, the overly optimistic $Static_{exp}$ increases the overall utility at the cost of a few days with zero utility. Note that the system achieved 100% uptime, but not always a sampling time of at least 2 hours, thus resulting in zero utility. As discussed in *Section 3.5*, $LT - ENO$ and $LT - ENO^+$ trade-off performance for ensuring continuous operation. When compared to statically setting the performance level, significantly higher utility can be achieved with a dynamic scheme due to the system's ability to react to deviations from expected conditions, and the resulting adjustment of the performance to safe levels.

**Fig. 6.4:**

Uninterrupted operation over 700 days for the energy harvesting system at location GG.

(top) $E_{in}(d)$ refers to the offline model, which, in this case, consistently overestimated the true conditions. In general, $E_{in}(d)$ can be an optimistic or pessimistic approximation, however. $E_{out}(t)$ refers to the expected energy consumption in Wh , while $Static_{exp}$ is the corresponding duty-cycle.

(bottom) $Static_{max}$ is the truly supported duty-cycle, assuming the actual conditions $E_{real}(t)$ are known at design-time. $LT - ENO^+$ and $LT - ENO$ refer to the duty-cycle achieved by the dynamic power management scheme from Chapter 3 with and without the modification discussed in Section 6.3.2. $SOC(t)$ refers to the battery State-of-Charge, i.e., the battery's fill-level. Note the zero input conditions around days 105 and 690, which trigger the emergency mechanism discussed in Section 3.3.2. This mechanism overrides the duty-cycle computed by $LT - ENO$ and $LT - ENO^+$ in an attempt to prevent low-voltage disconnects.




Fig. 6.5: Histogram of utility for the static ($Static_{max}$ and $Static_{exp}$), and dynamic approaches ($LT - ENO$ and $LT - ENO^+$).

Tab. 6.2: Percent Offline, mean (\overline{DC}) and minimum ($\lfloor DC \rfloor$) duty-cycle, duty-cycle variance normalized by the energy input variance (σ^2), mean utility (\overline{U}), minimum utility ($\lfloor U \rfloor$), and energy efficiency (η) for DH and GG shown in Figures 6.3 and 6.4. Note: $LT - ENO^+$ and $LT - ENO$ refer to the dynamic power management algorithms with and without the modification discussed in Section 6.3.2, while $Static_{exp}$ and $Static_{max}$ are expected static duty-cycle according to model assumptions.

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

6.4.1.2 OPT Performance Results

In the previous section we evaluated the dynamic power management algorithm from Chapter 3. To the best of our knowledge, to date, this is the only algorithm that is specifically designed for solar energy harvesting systems to achieve long-term uninterrupted operation at a stable




Fig. 6.6: (top) Energy use functions computed by the optimal ($u_{OPT}(t)$), clairvoyant ($u_{CV}(t)$), and LT-ENO ($u_{LT-ENO}(t)$) algorithms for the DH location. (bottom) The corresponding stored energy $SOC_x(t)$. Note that the time axis is given in weeks.

and deterministic minimum utilization. Therefore, it presents the perfect baseline for the evaluation of the optimal algorithm from *Chapter 4*.

The top graphs in *Figures 6.6 and 6.7* show the energy utilization functions for the DH and GG locations, computed with the baseline LT-ENO $u_{LT-ENO}(t)$, the optimal explicit algorithm $u_{OPT}(t)$ from *Section 4.3*, and, for reference, the clairvoyant optimal approach $u_{CV}(t)$ from *Section 4.4*, along with the input energy $p(t)$ and the corresponding battery states on the bottom graphs. Recall from *Section 4.2.1* that the energy utilization $u_x(t)$ is defined in terms of Watt-hours.

We first note that both algorithms exhibits remarkably low utilization variance that is maintained throughout the year. For the *LT – ENO* approach, this stems from the fact that the algorithm takes a long-term approach to compute the utilization by considering both the period of deficit, i.e., when the generation is below consumption, and the period of surplus, i.e., generation exceeds consumption. This approach attempts to fully leverage the battery to bridge periods of deficit, while simultaneously ensuring that the employed panel can actually recharge the battery during periods of surplus. The *OPT* approach, on the other hand, while leveraging the same long-term energy estimation model, computes the sustainable utilization by explicitly considering the battery state, which allows it to make much better use of surplus energy during summer.

The conservative approach by *LT – ENO* not only affects the minimum energy utilization, but also has a very detrimental effect on the total achievable system utility. From the battery profile shown on the bottom axes in *Figures 6.6 and 6.7*, it is obvious that the battery is full for most of the summer, yet the algorithm leverages neither the battery nor surplus energy by increasing the utilization. This stems from a fundamentally different approach, which attempts to ensure that the battery is at full capacity at the beginning of winter, at the cost of very conservative utilization during summer, and therefore reduced total system utility. In comparison, the *OPT* algorithm is able to leverage the surplus energy during summer, while ensuring that the battery is full at the appropriate time, and so ensure continuous operation and maximized total system utility.

Finally, we note that the minimum service-levels achieved by *LT – ENO* and the *OPT* approach are very similar during winter. In fact, during the winter periods, no clear winner can be established. However, when considering the entire period shown, the *OPT* algorithm achieves roughly 66% higher minimum utility, and 19% higher average energy use than *LT – ENO* for the DH location. For the GG location, *OPT* improves on *LT – ENO*'s average utility by 14%. The minimum utility is equivalent for both approaches, i.e., zero, which is due to extended periods of zero energy input, as discussed in the previous section. The worse performance by




Fig. 6.7: (top) Energy use functions computed by the optimal ($u_{OPT}(t)$), clairvoyant ($u_{CV}(t)$), and LT-ENO ($u_{LT-ENO}(t)$) algorithms for the GG location. (bottom) The corresponding stored energy $SOC_x(t)$. Note that the time axis is given in weeks.

Tab. 6.3: Percent Offline, minimum *weekly* energy utilization ($\lfloor u \rfloor$), total utilization ($\sum u$), utilization variance normalized by the energy input variance² (σ^2), average (\bar{U}) and minimum utility ($\lfloor U \rfloor$), and energy efficiency (η) for DH and GG shown in Figures 6.6 and 6.7.

Note: CV refers to the optimal clairvoyant algorithm from Section 4.3, OPT refers to the explicit algorithm from Section 4.4, while LT – ENO refers to the algorithm from Chapter 3.

	Power Management Algorithm	Offline [%]	$\lfloor u \rfloor$	$\sum u$	σ^2 [%]	\bar{U} [%]	$\lfloor U \rfloor$	η [%]
DH	CV	0	137.6	46,137	0.7088	90.5	59.1	100
	OPT	0	108.5	30,644	0.1943	82.9	37.3	66.42
	LT – ENO	0	84.0	20,176	1.7004	69.6	22.4	45.31
GG	CV	0	123.3	46,122	0.6774	91.0	47.2	100
	OPT	0	75.5	23,175	0.1347	80.1	0	50.24
	LT – ENO	0	18.2	21,220	2.3613	69.7	0	45.36

LT – ENO is mainly due to its pessimistic behavior in summer, where the use function tends to be lowest, despite the energy surplus that keeps the battery at full capacity.

In summary, the comparable utilization during winter can be mainly attributed to the fact that both algorithms use a very similar estimation model, and are able to leverage the available battery capacity. LT – ENO’s very pessimistic approach that fails to leverage surplus energy during summer, significantly hurts its overall performance. The service-level achieved by the OPT algorithm, on the other hand, is constrained only by the approximation accuracy of the estimator $\tilde{p}(t)$, and can maximize both the minimum utilization and utility simultaneously.

6.5 Chapter Summary

In this chapter we have shown that both, the LT – ENO approach, which is based on a heuristic, and the OPT approach, which has been shown in Chapter 4 to be optimal, are able to adapt to sub-optimally provisioned power subsystems and varying energy conditions, and still achieve uninterrupted long-term operation. Using the X-Sense [BBF⁺11] project as a concrete case study, we have exemplified how the lack of appropriate sensors can be circumvented, and so demonstrated our dynamic power

management schemes' applicability to solar energy harvesting systems with limited hardware support. Moreover, with evaluation over 700 days, we have shown that long-term uninterrupted operation of solar energy harvesting systems is indeed possible, and that the proposed approaches result in significant improvements in system utility and energy efficiency.

It is worth noting at this point that the work described in this thesis has enabled a number of interdisciplinary publications, a selection of which is given below:

- Samuel Weber, Stephan Gruber, Lucas Girard, and Jan Beutel. **Design of a measurement assembly to study in-situ rock damage driven by freezing.** In Proceeding of the 10th Int. Conference on Permafrost, Salekhard, Russia, 2012.
- Vanessa Wirz, Jan Beutel, Bernhard Buchli, Stephan Gruber, and Philippe Limpach. **Temporal characteristics of different cryosphere-related slope movements in high mountains.** In *Landslide Science and Practice*, pages 383–390. Springer, 2013.
- Zhenzhong Su, Alain Geiger, Philippe Limpach, Jan Beutel, Tonio Gsell, Bernhard Buchli, Stephan Gruber, Vanessa Wirz, and Felix Sutton. **Online monitoring of alpine slope instabilities with L1 GPS Real Time Kinematic Positions.** In EGU General Assembly Conference Abstracts, volume 16, page 7733, 2014.
- Vanessa Wirz, Marten Geertsema, Stephan Gruber, and Ross S Purves. **Temporal variability of diverse mountain permafrost slope movements derived from multi-year daily GPS data, Mattertal, Switzerland.** *Landslides*, pages 1–17, 2015.

7

Conclusions

We conclude this work with a brief review of the main results in *Section 7.1*, and a short discussion of possible further work in *Section 7.2*.

7.1 Main Results

This dissertation presented novel techniques for the design of the power subsystem and runtime power management for energy harvesting networked embedded systems. While our methods are generally applicable to a wide range of systems that leverage regenerative energy sources, throughout this thesis, we focused on solar powered Wireless Sensor Networks as the primary application scenario. We discussed how the proposed end-to-end power management solution enables uninterrupted system operation over multiple years despite relying on a highly variable energy source, the sun.

In this thesis we were not concerned with intricacies of digital system design, but focus on (i) how to appropriately provision the power subsystem (i.e., energy harvester and storage element), and (ii) light-weight approaches to dynamically adapting the system performance at runtime. Instead of attempting to reduce the system's energy consumption, subject to given performance constraints, our approach's goal is to maximize the system performance given the time-varying energy availability constraints. Compared to the State-of-the-Art methods, our solution achieves better, or at worst, equivalent performance but at a reduced system cost, i.e., requiring a smaller solar panel and/or battery.

In order to enable the system designer to adequately provision the power subsystem, we presented a capacity planning algorithm that leverages an energy availability model based on an astronomical radiation model. We showed through measurements with a solar energy harvesting setup that the proposed energy availability model can very accurately predict the long-term average harvestable energy. Based on this model, we presented an algorithm to compute the minimum battery and panel combination, given installation and technology parameters. With extensive simulation with input data for 10 years at 3 different geographical locations, we illustrated that an appropriately provisioned power subsystem is indispensable for achieving uninterrupted long-term operation.

A static, i.e., design-time approach, however, may not suffice to achieve uninterrupted operation if the conditions observed at the deployment site deviate too significantly from the design-time expectations. Therefore, to ensure that the system can efficiently use the available energy and operate without interruption over long time periods even in the presence of large deviations from the expectations, we presented two methods to dynamically adapt the system consumption at runtime. Through extensive simulation, we showed that our approaches enable uninterrupted operation at a pre-defined, stable minimum performance level over multiple years, despite the highly variable energy input observed. Moreover, we showed both analytically and through simulation that the proposed optimal approach maximizes both minimum and total energy utilization, dependent only on the accuracy of the energy estimation model. In comparison to three State-of-the-Art power management algorithms, our algorithms significantly improve the system's energy efficiency and so achieve improved system utility without experiencing system outages due to depleted batteries.

Finally, to further illustrate the benefits of our end-to-end solution, we evaluated it in the context of a case study based on a real-world scientific project. Since the deployed systems lack appropriate hardware support, we showed how a light-weight battery State-of-Charge algorithm can mitigate the systems' limitations, and so provide the dynamic power management algorithm with the necessary runtime information. Using two years of data collected with the deployed *X-Sense* wireless sensing system, we demonstrated the significant improvements in system utility that can be achieved with appropriate power subsystem design in combination with runtime management.

In summary, the main results of this thesis are as follows.

- We proposed a novel power subsystem capacity planning algorithm that leverages an energy availability model based on an astronomical solar radiation model. Through measurements with a solar

panel, we validated that the energy availability model, leveraged by the capacity planning algorithm and the dynamic power management schemes, performs very well at approximating the long-term average harvestable energy. We further showed that our approach outperforms the State-of-the-Art in capacity planning by achieving 100% system availability over ten years of simulation, while requiring up to 53% smaller batteries for the evaluated datasets.

- We presented two new dynamic power management algorithms that maximize the minimum, and total energy utilization, i.e., service-level, with very little variance, i.e., high service-level stability. Through extensive simulation we show that our algorithms significantly outperform the State-of-the-Art techniques in the applied metrics, while incurring zero downtime. In fact, the proposed optimal algorithm achieves to within 9.9% of the theoretical optimal performance. We presented light-weight implementation for both algorithms, suitable for realization on resource-constrained systems.
- We introduced a practical, trace-based, direct-measurement method for online battery State-of-Charge approximation. Through experimentation we showed that acceptable approximation accuracies are possible even for situations that are not explicitly modeled, e.g., aging and temperature effects. We further exemplified how the algorithm can be leveraged to mitigate lack of adequate hardware measurement support.
- Finally, using a concrete, real-world project as a case study, we demonstrated that our end-to-end power management solution can significantly improve the end-user application utility, without requiring extensive information about the energy conditions at the intended deployment site.

7.2 Possible Future Perspectives

The work discussed in this thesis presents fundamental insights into the challenges in design and operation of solar energy harvesting systems. While we believe that our approach has the potential to advance solar energy harvesting systems from opportunistic consumers of “surplus” energy to systems that efficiently, and most importantly, safely leverage available energy to achieve dependable long-term operation, we can identify a number of topics that deserve further investigation. In the following we briefly identify avenues for potential future work for each of the *Chapters* 2 through 5.

Chapter 2: Power Subsystem Capacity Planning

From the discussion in *Chapter 2*, two possible improvements are evident. First, considering the highly non-linear behavior in energy conversion and particularly storage, incorporation of these effects into the energy availability and battery models may be beneficial. Specifically, since we aim at multi-season operation, possibly large temperature variations may be observed. As explained in *Chapter 5*, reduced temperatures result in lower apparent battery capacity. Since periods of lower temperatures coincide with periods of reduced solar radiation, accounting for the reduced storage capacity may be critical in achieving uninterrupted operation. Moreover, since batteries and solar panels tend to degrade over time, the incorporation of degradation factors may further ensure that the computed power subsystem can actually support the expected performance over multiple years. Recall that the capacity planning is performed offline, hence the incurred complexity due to more advanced models is unlikely to be an issue.

Chapters 3 and 4: Dynamic Power Management

Incorporation of component degradation due to age and especially temperature, as described above, is one of the immediate concerns for the proposed dynamic power management schemes. In fact, considering that the uncertainty in energy predictions is the main contributor deviations from the optimal, it is intuitively clear that our proposed power management schemes ultimately benefit from improved energy prediction accuracy. In fact, for the evaluated datasets, the OPT approach from *Chapter 4* comes to within 9.9% of the theoretical possible, i.e., with perfect prediction accuracy. This illustrates that one of the most promising directions for future work lies in the investigation of ways to improve the energy prediction accuracy.

Chapter 5: Battery State-of-Charge

As already discussed in *Chapter 5*, the charging model and the finite state machine to track the individual states present room for improvement. While the proposed approach works well for system with limited measurement support, a combination of coulomb counter and software house-keeping is certainly the preferred method for State-of-Charge approximation. In fact, for certain battery chemistries, e.g., Lithium-ion or Lithium-polymer batteries, the proposed approach would not work due to the stable terminal voltage exhibited by these batteries.

Distributed Applications and/or Different Harvesting Sources

The power management schemes discussed in this work are concerned with local, i.e., per mote, optimization, and we have neglected network wide optimization. However, in many application scenarios the motes jointly maximize the application objective. In this case, the spatial variations in environmental energy availability must be exploited. Moreover, the individual systems may be powered by different sources, which may exhibit very different energy availability profiles. Therefore, the probably most interesting venue is the extension of our work to distributed application scenarios, possibly coupled with different harvesting sources. Achieving multi-source, network-wide optimization would certainly be an interesting and challenging problem to solve. Finally, it may be worth the effort to investigate if, and how the techniques from *Chapters 3 and 4* apply to systems that lack energy storage elements and can therefore only operate when energy is being harvested from the environment.

Bibliography

- [AAA⁺07] C. Arnhardt, K. Asch, R. Azzam, R. Bill, TM. Fernandez-Steeger, SD. Homfeld, A. Kallash, F. Niemeyer, H. Ritter, M. Toloczyki, et al. Sensor based Landslide Early Warning System-SLEWS. Development of a geoservice infrastructure as basis for early warning systems for landslides by integration of real-time sensors. *Geotechnologien science report*, 10:75–88, 2007.
- [AAP⁺12] Daniele Alessandrelli, Andrea Azzarà, Matteo Petracca, Christian Nastasi, and Paolo Pagano. ScanTraffic: Smart Camera Network for Traffic Information Collection. In *Wireless Sensor Networks*, pages 196–211. Springer Science & Business Media, 2012.
- [AC92] J. Scott Armstrong and Fred Collopy. Error measures for generalizing about forecasting methods: Empirical comparisons. *International Journal of Forecasting*, 8(1):69–80, Jun 1992.
- [AHS07] Karl Aberer, Manfred Hauswirth, and Ali Salehi. Infrastructure for data processing in large-scale interconnected sensor networks. In *Mobile Data Management, 2007 International Conference on*, pages 198–205. IEEE, 2007.
- [AUDT10] Vivek Agarwal, Kasemsak Uthaichana, Raymond A DeCarlo, and Lefteri H Tsoukalas. Development and validation of a battery model useful for discharging and charging power control and lifetime estimation. *Energy Conversion, IEEE Transactions on*, 25(3):821–835, 2010.
- [BAB13] Bernhard Buchli, Daniel Aschwanden, and Jan Beutel. Battery State-of-Charge Approximation for Energy Harvesting Embedded Systems. In *Wireless Sensor Networks*, pages 179–196. Springer Science & Business Media, 2013.

- [BBB09] Carlo Bergonzini, Davide Brunelli, and Luca Benini. Algorithms for harvested energy prediction in batteryless wireless sensor networks. In *2009 3rd International Workshop on Advances in sensors and Interfaces*, pages 144–149. IEEE, Jun 2009.
- [BBF⁺11] Jan Beutel, Bernhard Buchli, Federico Ferrari, Matthias Keller, Marco Zimmerling, and Lothar Thiele. X-Sense: Sensing in extreme environments. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2011*, pages 1–6. IEEE, 2011.
- [BBMT08] Davide Brunelli, Luca Benini, Clemens Moser, and Lothar Thiele. An efficient solar energy harvester for wireless sensor nodes. In *Proceedings of the conference on Design, automation and test in Europe*, pages 104–109. ACM, Mar 2008.
- [BC06] Craig F Bohren and Eugene E Clothiaux. *Fundamentals of atmospheric radiation: an introduction with 400 problems*. John Wiley & Sons, 2006.
- [BGH⁺09] Jan Beutel, Stephan Gruber, Andreas Hasler, Roman Lim, Andreas Meier, Christian Plessl, Igor Talzi, Lothar Thiele, Christian Tschudin, Matthias Woehrle, et al. PermaDAQ: A scientific instrument for precision sensing and data recovery in environmental extremes. In *Proceedings of the 2009 International Conference on Information Processing in Sensor Networks*, pages 265–276. IEEE Computer Society, 2009.
- [BKN02] Henk Jan Bergveld, Wanda S Kruijt, and Peter HL Notten. *Battery management systems*. Springer, 2002.
- [BKT15] Bernhard Buchli, Pratyush Kumar, and Lothar Thiele. Optimal Power Management With Guaranteed Minimum Energy Utilization For Solar Energy Harvesting Systems. In *Distributed Computing in Sensor Systems (DCOSS), 2015 IEEE International Conference on*. IEEE, 2015.
- [BMN09] Peder Bacher, Henrik Madsen, and Henrik Aalborg Nielsen. Online short-term solar power forecasting. *Solar Energy*, 83(10):1772–1783, 2009.
- [BPBB12] D. Balsamo, G. Paci, L. Benini, and D. Brunelli. Temperature compensated time synchronisation in wireless sensor networks. *Electronics Letters*, 48(16):1026–1028, aug 2012.

- [BRVK11] Ravneet Bajwa, Ram Rajagopal, Pravin Varaiya, and Robert Kavalier. In-pavement wireless sensor network for vehicle classification. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 85–96. IEEE, 2011.
- [BSB12] Bernhard Buchli, Felix Sutton, and Jan Beutel. GPS-Equipped Wireless Sensor Network Node for High-Accuracy Positioning Applications. In *Wireless Sensor Networks*, pages 179–195. Springer Science & Business Media, 2012.
- [BSBT14a] Bernhard Buchli, Felix Sutton, Jan Beutel, and Lothar Thiele. Dynamic power management for long-term energy neutral operation of solar energy harvesting systems. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems - SenSys '14*, pages 31–45. ACM Press, 2014.
- [BSBT14b] Bernhard Buchli, Felix Sutton, Jan Beutel, and Lothar Thiele. Towards Enabling Uninterrupted Long-Term Operation of Solar Energy Harvesting Embedded Systems. In *Wireless Sensor Networks*, pages 66–83. Springer Science & Business Media, 2014.
- [Buz08] Matteo Buzzi. *Challenges in operational numerical weather prediction at high resolution in complex terrain*. PhD thesis, Diss., Eidgenössische Technische Hochschule ETH Zürich, Nr. 17714, 2008.
- [BVW07] Nicolas Burri, Pascal Von Rickenbach, and Roger Wattenhofer. Dozer: ultra-low power data gathering in sensor networks. In *Information Processing in Sensor Networks, 2007. IPSN 2007. 6th International Symposium on*, pages 450–459. IEEE, 2007.
- [BYL⁺11] Bernhard Buchli, Mustafa Yücel, Roman Lim, Tonio Gsell, and Jan Beutel. Demo abstract: Feature-rich platform for WSN design space exploration. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 115–116, April 2011.
- [CB10] Marcus Chang and Philippe Bonnet. Meeting ecologists' requirements with adaptive data acquisition. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 141–154. ACM, 2010.

- [CC08] Sravanthi Chalasani and James M Conrad. A survey of energy harvesting sources for embedded systems. In *South-eastcon, 2008. IEEE*, pages 442–447. IEEE, 2008.
- [CD05] Mihaela Cardei and Ding-Zhu Du. Improving wireless sensor network lifetime through power aware organization. *Wireless Networks*, 11(3):333–340, 2005.
- [cle] cleversolar. 30W Solar Panel Professional Monocrystalline CS-30. [Online; accessed 08-Mar-2015], <http://www.cleversolar.co.uk/shop/solar-panels/monocrystalline-solar-panels/30w-solar-panel-professional-982978.html>.
- [Con11] Concorde Battery Corporation. *Technical Manual for Lifeline Batteries*, July 2011. [Online; accessed 08-Mar-2015], <http://www.lifelinebatteries.com/manual.pdf>.
- [CPBA12] Andrea Castagnetti, Alain Pegatoquet, Cécile Belleudy, and Michel Auguin. An efficient state of charge prediction model for solar harvesting wsn platforms. In *Systems, Signals and Image Processing (IWSSIP), 2012 19th International Conference on*, pages 122–125. IEEE, 2012.
- [CPS12] Alessandro Cammarano, Chiara Petrioli, and Dora Spenza. Pro-Energy: a novel energy prediction model for solar and wind energy-harvesting Wireless Sensor Networks. *Proceedings of IEEE MASS 2012*, pages 8–11, 2012.
- [CSSJ11] Shengbo Chen, Prasun Sinha, Ness B Shroff, and Changhee Joo. Finite-horizon energy allocation and routing scheme in rechargeable sensor networks. In *INFOCOM, 2011 Proceedings IEEE*, pages 2273–2281. IEEE, 2011.
- [CVS⁺07] Peter Corke, Philip Valencia, Pavan Sikka, Tim Wark, and Les Overs. Long-duration solar-powered wireless sensor networks. In *Proceedings of the 4th workshop on Embedded networked sensors, EmNets '07*, pages 33–37, New York, NY, USA, 2007. ACM.
- [DAS06] Dennis Doerffel and Suleiman Abu-Sharkh. A critical review of using Peukert-equation for determining the remaining capacity of lead-acid and lithium ion batteries. *Journal of Power Sources*, 155(2):395–400, 2006.

- [DEM⁺10] Vladimir Dyo, Stephen A Ellwood, David W Macdonald, Andrew Markham, Cecilia Mascolo, Bence Pásztor, Salvatore Scellato, Niki Trigoni, Ricklef Wohlers, and Kharsim Yousef. Evolution and sustainability of a wildlife monitoring sensor network. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 127–140. ACM, 2010.
- [DFFMM06] Henri Dubois-Ferrière, Laurent Fabre, Roger Meier, and Pierre Metrailler. TinyNode: a comprehensive platform for wireless sensor network applications. In *Proceedings of the 5th international conference on Information processing in sensor networks*, pages 358–365. ACM, 2006.
- [DHM75] J.V. Dave, P. Halpern, and H.J. Myers. Computation of Incident Solar Energy. *IBM Journal of Research and Development*, 19(6):539–549, 1975.
- [DOTH07] Adam Dunkels, Fredrik Osterlind, Nicolas Tsiftes, and Zhi-tao He. Software-based on-line energy estimation for sensor nodes. In *Proceedings of the 4th workshop on Embedded networked sensors*, pages 28–32. ACM, 2007.
- [FGJK13] Tat S Fu, Amitabha Ghosh, Erik A Johnson, and Bhaskar Krishnamachari. Energy-efficient deployment strategies in structural health monitoring using wireless sensor networks. *Structural Control and Health Monitoring*, 20(6):971–986, 2013.
- [FZTS11] Federico Ferrari, Marco Zimmerling, Lothar Thiele, and Olga Saukh. Efficient network flooding and time synchronization with Glossy. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 73–84. IEEE, 2011.
- [GEH⁺12] Martin A Green, Keith Emery, Yoshihiro Hishikawa, Wilhelm Warta, and Ewan D Dunlop. Solar cell efficiency tables (version 39). *Progress in photovoltaics: research and applications*, 20(1):12–20, 2012.
- [GGP12] S. Gubler, S. Gruber, and R.S. Purves. Uncertainties of parameterized surface downward clear-sky shortwave and all-sky longwave radiation. *Atmospheric Chemistry and Physics*, 12(11):5077–5098, 2012.

- [GHL⁺87] Leonidas Guibas, John Hershberger, Daniel Leven, Micha Sharir, and Robert E Tarjan. Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*, 2(1-4):209–233, 1987.
- [GHSW10] Philipp Maria Glatz, Leander Bernd Hörmann, Christian Steger, and Reinhold Weiss. 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*, pages 55–65, 2010.
- [GSGSGH11] Antonio-Javier Garcia-Sanchez, Felipe Garcia-Sanchez, and Joan Garcia-Haro. Wireless sensor network deployment for integrating video-surveillance and data-monitoring in precision agriculture over distributed crops. *Computers and Electronics in Agriculture*, 75(2):288–303, 2011.
- [gum] Gumstix. [Online; accessed 08-Mar-2015], <http://www.gumstix.com>.
- [HG10] L Hanssen and J Gakkestad. Solar Cell Size Requirement for Powering of Wireless Sensor Network Used in Northern Europe. In *Proceedings of the International Workshops on PowerMEMS*, pages 17–20, 2010.
- [HKS09] Jaehyun Han, Dongchul Kim, and Myoungho Sunwoo. State-of-charge estimation of lead-acid batteries using an adaptive extended Kalman filter. *Journal of Power Sources*, 188(2):606–612, 2009.
- [HLG06] Detlev Heinemann, Elke Lorenz, and Marco Girodo. Forecasting of solar radiation. *Solar energy resource management for electricity generation from local level to global scale*. Nova Science Publishers, New York, 2006.
- [HTB⁺08] Andreas Hasler, Igor Talzi, Jan Beutel, Christian Tschudin, and Stephan Gruber. Wireless sensor networks in permafrost research-concept, requirements, implementation and challenges. In *Proc. 9th Int’l Conf. on Permafrost (NICOP 2008)*, volume 1, pages 669–674, 2008.
- [Hun86] J. Stuart Hunter. The exponentially weighted moving average. *J. QUALITY TECHNOL.*, 18(4):203–210, 1986.

- [HVV⁺06] Tian He, Pascal Vicaire, Ting Yan, Liqian Luo, Lin Gu, Gang Zhou, Radu Stoleru, Qing Cao, John A Stankovic, and Tarek Abdelzaher. Achieving real-time target tracking using wireless sensor networks. In *Real-Time and Embedded Technology and Applications Symposium, 2006. Proceedings of the 12th IEEE*, pages 37–48. IEEE, 2006.
- [Jam08] Farooq Balouchi James M. Gilbert. Comparison of Energy Harvesting Systems for Wireless Sensor Networks. *International Journal of Automation and Computing*, 5(4):334, 2008.
- [JC12] Jaein Jeong and David Culler. Predicting the Long-Term Behavior of a Micro-Solar Power System. *ACM Transactions on Embedded Computing Systems (TECS)*, 11(2):35, 2012.
- [JCDS10] Raja Jurdak, Peter Corke, Dhinesh Dharman, and Guillaume Salagnac. Adaptive GPS duty cycling and radio ranging for energy-efficient localization. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 57–70. ACM, 2010.
- [JDCS07] Xiaofan Jiang, Prabal Dutta, David Culler, and Ion Stoica. Micro power meter for energy monitoring of wireless sensor networks at scale. In *Proceedings of the 6th international conference on Information processing in sensor networks*, pages 186–195. ACM, 2007.
- [KDMB12] Hessam Kooti, Nga Dang, Deepak Mishra, and Eli Bozorgzadeh. Energy Budget Management for Energy Harvesting Embedded Systems. In *Embedded and Real-Time Computing Systems and Applications (RTCSA), 2012 IEEE 18th International Conference on*, pages 320–329. IEEE, 2012.
- [KH06] Wook Hyun Kwon and Soo Hee Han. *Receding horizon control: model predictive control for state models*. Springer, 2006.
- [KHSR06] Aman Kansal, Jason Hsu, Mani Srivastava, and Vijay Raghunathan. Harvesting aware power management for sensor networks. In *Proceedings of the 43rd annual Design Automation Conference*, pages 651–656. ACM, 2006.
- [KHZS07] Aman Kansal, Jason Hsu, Sadaf Zahedi, and Mani B Srivastava. Power management in energy harvesting sensor networks. *ACM Transactions on Embedded Computing Systems (TECS)*, 6(4):32, 2007.

- [KLS⁺10] JeongGil Ko, Chenyang Lu, Mani B Srivastava, John A Stankovic, Andreas Terzis, and Matt Welsh. Wireless sensor networks for healthcare. *Proceedings of the IEEE*, 98(11):1947–1960, 2010.
- [KPA⁺10] Fotis Kerasiotis, Aggeliki Prayati, Christos Antonopoulos, Christos Koulamas, and George Papadopoulos. Battery Lifetime Prediction Model for a WSN Platform. In *Proceedings of the 2010 Fourth International Conference on Sensor Technologies and Applications, SENSORCOMM '10*, pages 525–530, Washington, DC, USA, 2010. IEEE Computer Society.
- [KPS04] Aman Kansal, Dunny Potter, and Mani B Srivastava. Performance aware tasking for environmentally powered sensor networks. In *ACM SIGMETRICS Performance Evaluation Review*, volume 32, pages 223–234. ACM, 2004.
- [KS03] A. Kansal and M.B. Srivastava. An environmental energy harvesting framework for sensor networks. In *Low Power Electronics and Design, 2003. ISLPED '03. Proceedings of the 2003 International Symposium on*, pages 481–486, 2003.
- [LGB⁺11] Philippe Limpach, Alain Geiger, Jan Beutel, Bernhard Buchli, Vanessa Wirz, and Stephan Gruber. Permanent monitoring of rock glaciers with low-cost GPS. In *9th Swiss Geoscience Meeting*, pages 373–373, 2011.
- [Lim] Dr. Philippe Limpach. personal communication.
- [LK07] Fajie Li and Reinhard Klette. Rubberband algorithms for solving various 2D or 3D shortest path problems. In *Computing: Theory and Applications, 2007. ICCTA'07. International Conference on*, pages 9–19. IEEE, 2007.
- [LK11] Fajie Li and Reinhard Klette. *Euclidean Shortest Paths*. Springer, 2011.
- [LMP⁺05] Philip Levis, Sam Madden, Joseph Polastre, Robert Szewczyk, Kamin Whitehouse, Alec Woo, David Gay, Jason Hill, Matt Welsh, Eric Brewer, et al. Tinyos: An operating system for sensor networks. In *Ambient intelligence*, pages 115–148. Springer, 2005.
- [LSB⁺12] Trong Nhan Le, Olivier Sentieys, Olivier Berder, Alain Pegatoquet, and Cecile Belleudy. Power Manager with PID

- controller in Energy Harvesting Wireless Sensor Networks. In *Green Computing and Communications (GreenCom), 2012 IEEE International Conference on*, pages 668–670. IEEE, 2012.
- [LW12] Jiakang Lu and Kamin Whitehouse. SunCast: fine-grained prediction of natural sunlight levels for improved daylight harvesting. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, pages 245–256. ACM, 2012.
- [Max98] Maxim Integrated Products. Evaluation Kit for the MAX1660. <http://datasheets.maxim-ic.com/en/ds/MAX1660EVKIT.pdf> [Accessed: March 8, 2015], 1998.
- [MBTB07] Clemens Moser, Davide Brunelli, Lothar Thiele, and Luca Benini. Real-time scheduling for energy harvesting sensor nodes. *Real-Time Systems*, 37(3):233–260, 2007.
- [ML09] Raquel A.F. Mini and Antonio A.F. Loureiro. Energy in wireless sensor networks. In *Middleware for Network Eccentric and Mobile Applications*, pages 3–24. Springer, 2009.
- [MOH04] K. Martinez, R. Ong, and J. Hart. Glacsweb: a sensor network for hostile environments. In *Sensor and Ad Hoc Communications and Networks, 2004. IEEE SECON 2004. 2004 First Annual IEEE Communications Society Conference on*, pages 81–87, October 2004. Doi: 10.1109/SAHCN.2004.1381905.
- [Mor09] Morningstar Corporation. *SunSaver Photovoltaic System Controllers Operator’s Manual*, November 2009. [Online; accessed 08-Mar-2015], <http://www.morningstarcorp.com/en/support/library/SS.IOM.01.EN.pdf>.
- [PAG09] Maria Teresa Penella, Joan Albasa, and Manel Gasulla. Powering wireless sensor nodes: Primary batteries versus energy harvesting. In *IEEE Instrumentation and Measurement Technology Conference*, 2009.
- [PBAR09] Joaquin Recas Piorno, Carlo Bergonzini, David Atienza, and Tajana Simunic Rosing. Prediction and management in energy harvested wireless sensor nodes. In *Wireless Communication, Vehicular Technology, Information Theory and*

- Aerospace & Electronic Systems Technology, 2009. Wireless VI-TAE 2009. 1st International Conference on*, pages 6–10. IEEE, 2009.
- [PBNR05] Valer Pop, Henk Jan Bergveld, PHL Notten, and Paul PL Regtien. State-of-the-art of battery state-of-charge determination. *Measurement Science and Technology*, 16(12):R93, 2005.
- [PC06] Chulsung Park and Pai H Chou. 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, pages 168–177. IEEE, 2006.
- [Peu97] W. Peukert. Über die Abhängigkeit der Kapazität von der Entladestromstärke bei Bleiakkumulatoren. *Elektrotechnische Zeitschrift*, 20:20–21, 1897.
- [PLR05] Chulsung Park, Kanishka Lahiri, and Anand Raghunathan. Battery discharge characteristics of wireless sensor nodes: An experimental analysis. In *In Proceedings of the IEEE Conf. on Sensor and Ad-hoc Communications and Networks (SECON)*, 2005.
- [PPJ01] Sabine Piller, Marion Perrin, and Andreas Jossen. Methods for state-of-charge determination and their applications. *Journal of Power Sources*, 96(1):113–120, 2001.
- [Ran06] Julian Randall. *Designing indoor solar products*. John Wiley & Sons, 2006.
- [RC06] Vijay Raghunathan and Pai H. Chou. Design and power management of energy harvesting embedded systems. In *Proceedings of the 2006 international symposium on Low power electronics and design, ISLPED'06*, pages 369–374, New York, NY, USA, 2006. ACM.
- [RGS06] Vijay Raghunathan, Saurabh Ganeriwal, and Mani Srivastava. Emerging techniques for long lived wireless sensor networks. *Communications Magazine, IEEE*, 44(4):108–114, 2006.
- [RKH⁺05] Vijay Raghunathan, Aman Kansal, Jason Hsu, Jonathan Friedman, and Mani Srivastava. Design considerations for solar energy harvesting wireless embedded systems. In

- Proceedings of the 4th international symposium on Information processing in sensor networks*, page 64. IEEE Press, 2005.
- [RVR03] Ravishankar Rao, Sarma Vrudhula, and Daler N Rakhmatov. Battery modeling for energy aware system design. *Computer*, 36(12):77–87, 2003.
- [SCO⁺07] P. Sitka, Peter Corke, Leslie Overs, Philip Valencia, and Tim Wark. Fleck-a platform for real-world outdoor sensor networks. In *Intelligent Sensors, Sensor Networks and Information, 2007. ISSNIP 2007. 3rd International Conference on*, pages 709–714. IEEE, 2007.
- [SGIS10] Navin Sharma, Jeremy Gummeson, David Irwin, and Prashant Shenoy. Cloudy computing: Leveraging weather forecasts in energy harvesting sensor systems. In *Sensor Mesh and Ad Hoc Communications and Networks (SECON), 2010 7th Annual IEEE Communications Society Conference on*, pages 1–9. IEEE, 2010.
- [SGL⁺14] Zhenzhong Su, Alain Geiger, Philippe Limpach, Jan Beutel, Tonio Gsell, Bernhard Buchli, Stephan Gruber, Vanessa Wirz, and Felix Sutton. Online monitoring of alpine slope instabilities with L1 GPS Real Time Kinematic Positions. In *EGU General Assembly Conference Abstracts*, volume 16, page 7733, 2014.
- [SK11] Sujesha Sudevalayam and Purushottam Kulkarni. Energy harvesting sensor nodes: Survey and implications. *Communications Surveys & Tutorials, IEEE*, 13(3):443–461, 2011.
- [SKJ13] Philipp Sommer, Branislav Kusy, and Raja Jurdak. Power management for long-term sensing applications with energy harvesting. In *Proceedings of the 1st International Workshop on Energy Neutral Sensing Systems*, page 3. ACM, 2013.
- [SKMJ14] Philipp Sommer, Branislav Kusy, Adam McKeown, and Raja Jurdak. The big night out: Experiences from tracking flying foxes with delay-tolerant wireless networking. In *Real-World Wireless Sensor Networks*, pages 15–27. Springer, 2014.
- [SSIS11] Navin Sharma, Pranshu Sharma, David Irwin, and Prashant Shenoy. Predicting solar generation from weather

- forecasts using machine learning. In *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on*, pages 528–533. IEEE, 2011.
- [STC13] Winston K.G. Seah, Y.K. Tan, and Alvin T.S. Chan. *Research in energy harvesting wireless sensor networks and the challenges ahead*. Springer, 2013.
- [Ste13] Mathis Steichen. NiMH Battery State-of-Charge Approximation. *Semester Thesis, ETH Zurich*, June 2013. [Online; accessed 08-Mar-2015], <http://ftp.tik.ee.ethz.ch/pub/students/2013-FS/SA-2013-05.pdf>.
- [TJC08] Jay Taneja, Jaemin Jeong, and David Culler. Design, Modeling, and Capacity Planning for Micro-solar Power Sensor Networks. In *Proc. of the 7th international conference on Information processing in sensor networks, IPSN '08*, pages 407–418, Washington, DC, USA, 2008. IEEE Computer Society.
- [VGB07] Christopher M Vigorito, Deepak Ganesan, and Andrew G Barto. Adaptive control of duty cycling in energy-harvesting wireless sensor networks. In *Sensor, Mesh and Ad Hoc Communications and Networks, 2007. SECON'07. 4th Annual IEEE Communications Society Conference on*, pages 21–30. IEEE, 2007.
- [VKR⁺05] Iuliu Vasilescu, Keith Kotay, Daniela Rus, Matthew Dumbabin, and Peter Corke. Data Collection, Storage, and Retrieval with an Underwater Sensor Network. In *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems, SenSys '05*, pages 154–165, New York, NY, USA, 2005. ACM.
- [VPB07] Amir Vasebi, Maral Partovibakhsh, and S Mohammad Taghi Bathaee. A novel combined battery model for state-of-charge estimation in lead-acid batteries based on extended Kalman filter for hybrid electric vehicle applications. *Journal of Power Sources*, 174(1):30–40, 2007. Hybrid Electric Vehicles.
- [VST13] Alvin C. Valera, Wee-Seng Soh, and Hwee-Pink Tan. Energy-neutral scheduling and forwarding in environmentally-powered wireless sensor networks. *Ad Hoc Networks*, 11(3):1202–1220, 2013.

- [WALR⁺06] Geoffrey Werner-Allen, Konrad Lorincz, Mario Ruiz, Omar Marcillo, Jeff Johnson, Jonathan Lees, and Matt Welsh. Deploying a wireless sensor network on an active volcano. *Internet Computing, IEEE*, 10(2):18–25, 2006.
- [WBB⁺13] Vanessa Wirz, Jan Beutel, Bernhard Buchli, Stephan Gruber, and Philippe Limpach. Temporal characteristics of different cryosphere-related slope movements in high mountains. In *Landslide Science and Practice*, pages 383–390. Springer, 2013.
- [WGGB12] Samuel Weber, Stephan Gruber, Lucas Girard, and Jan Beutel. Design of a measurement assembly to study in-situ rock damage driven by freezing. In *Proceeding of the 10th Int. Conference on Permafrost, Salekhard, Russia*, 2012.
- [WGGP15] Vanessa Wirz, Marten Geertsema, Stephan Gruber, and Ross S Purves. Temporal variability of diverse mountain permafrost slope movements derived from multi-year daily GPS data, Matternal, Switzerland. *Landslides*, pages 1–17, 2015.
- [Wip10] Jean-Claude Wippler. Jeenode. January 2010. [Online; accessed 01-Aug-2011], <http://jeelabs.org/2010/01/23/lithium-node/>.
- [WLLP01] Brett Warneke, Matt Last, Brian Liebowitz, and Kristofer SJ Pister. Smart dust: Communicating with a cubic-millimeter computer. *Computer*, 34(1):44–51, 2001.
- [ZRLM09] Kai Zeng, Kui Ren, Wenjing Lou, and Patrick J Moran. Energy aware efficient geographic routing in lossy wireless sensor networks with environmental energy supply. *Wireless Networks*, 15(1):39–51, 2009.
- [ZSLM04] Pei Zhang, Christopher M Sadler, Stephen A Lyon, and Margaret Martonosi. Hardware design experiences in ZebraNet. In *Proceedings of the 2nd international conference on Embedded networked sensor systems, SenSys '04*, pages 227–238, New York, NY, USA, 2004. ACM.

List of Publications

The following list includes publications that form the basis of this thesis. The corresponding chapters are indicated in parentheses.

(4,6) Bernhard Buchli, Pratyush Kumar, and Lothar Thiele. Optimal Power Management With Guaranteed Minimum Energy Utilization For Solar Energy Harvesting Systems. In *Distributed Computing in Sensor Systems (DCOSS), 2015 IEEE International Conference on*, Fortaleza, Brazil, June 2015. IEEE.

(3,6) Bernhard Buchli, Felix Sutton, Jan Beutel, and Lothar Thiele. Dynamic power management for long-term energy neutral operation of solar energy harvesting systems. In *Proceedings of the 12th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, Memphis, TN, USA, Nov 2014. ACM.

(2) Bernhard Buchli, Felix Sutton, Jan Beutel, and Lothar Thiele. Towards enabling uninterrupted long-term operation of solar energy harvesting embedded systems. In *Proceedings of 11th European Conference on Wireless Sensor Networks (EWSN 2014)*, Oxford, UK, Feb 2014. Springer Link, Lecture Notes on Computer Science.

(5) Bernhard Buchli, Daniel Aschwanden, and Jan Beutel. Battery state-of-charge approximation for energy harvesting embedded systems. In *Proceedings of 10th European Conference on Wireless Sensor Networks (EWSN 2013)*, pages 179–196, Ghent, Belgium, Feb 2013. Springer Link, Lecture Notes on Computer Science.

(6) Jan Beutel, Bernhard Buchli, Federico Ferrari, Matthias Keller, Lothar Thiele, and Marco Zimmerling. X-sense: Sensing in extreme environments. In *Proceedings of Design, Automation and Test in Europe (DATE 2011)*, pages 1460–1465, Grenoble, France, Mar 2011.

(6) Bernhard Buchli, Mustafa Yucel, Roman Lim, Tonio Gsell, and Jan Beutel. Demo abstract: Feature-rich experimentation for wsn design space exploration. In *Proceedings of the 10th International Conference on Information*

Processing in Sensor Networks (IPSN 2011), pages 115–116, Chicago, IL, USA, Apr 2011. ACM/IEEE.

The following list summarizes the publications that were written during the PhD studies, yet are not part of this thesis.

Samuel Weber, Jan Beutel, Bernhard Buchli, Stephan Gruber, Tonio Gsell, Roman Lim, Philippe Limpach, Hugo Raetzo, Zhenzhong Su, Felix Sutton, Christoph Walser, and Vanessa Wirz. Permasense l1-gps for kinematic monitoring. In *Book of Abstracts of EUCOP4 – 4 th European Conference on Permafrost*, pages 186–187, Évora, Portugal, Jun 2014. University of Lisbon and the University of Évora.

Zhenzhong Su, Alain Geiger, Philippe Limpach, Jan Beutel, Tonio Gsell, Bernhard Buchli, Stephan Gruber, Vanessa Wirz, , and Felix Sutton. Online monitoring of alpine slope instabilities with l1 gps real time kinematic positions. In *Geophysical Research Abstracts*, volume 16, page 7733, EGU General Assembly 2014, Apr 2014.

Vanessa Wirz, Jan Beutel, Bernhard Buchli, Stephan Gruber, and Philippe Limpach. Temporal characteristics of different cryosphere-related slope movements in high mountains. In *Landslide Science and Practice*, pages 383–390. Berlin, Heidelberg, Sep 2013.

Bernhard Buchli, Felix Sutton, and Jan Beutel. Gps-equipped wireless sensor network node for high-accuracy positioning applications. In *Proceedings of 9th European Conference on Wireless Sensor Networks (EWSN 2012)*, pages 179–195, Trento, Italy, Feb 2012. Springer Link, Lecture Notes on Computer Science.

Felix Sutton, Bernhard Buchli, and Jan Beutel. Demo abstract: An ultra-low power wireless multi-sensor interface. In *Lecture Notes on Computer Science. Proc. of 10th European Conference on Wireless Sensor Networks (EWSN 2013)*, EWSN 2013 Demo, Ghent, Belgium, Feb 2013. Springer.

Philippe Limpach, Alain Geiger, Jan Beutel, Bernhard Buchli, Vanessa Wirz, and Stephan Gruber. Permanent monitoring of rock glaciers with low-cost gps. In *Abstract Volume 9th Swiss Geoscience Meeting*, volume 9, page 373. Platform Geosciences, Swiss Academy of Science, SCNAT, Nov 2011.

Vanessa Wirz, Jan Beutel, Bernhard Buchli, Stephan Gruber, Philippe Limpach, and Zhenzhong Su. Temporal characteristics of different cryosphere-related slope movements in high mountains: Gps measurements and analysis. In *Abstract Volume 9th Swiss Geoscience Meeting*,

volume 9, page 282. Platform Geosciences, Swiss Academy of Science, SCNAT, Nov 2011.

Bernhard Buchli, Felix Sutton, and Jan Beutel. Gps-equipped wireless sensor network for geodetic positioning applications. In *Abstract Volume 9th Swiss Geoscience Meeting*, volume 17, pages 366–367. Platform Geosciences, Swiss Academy of Science, SCNAT, Nov 2011.

Vanessa Wirz, Philippe Limpach, Jan Beutel, Bernhard Buchli, and Stephan Gruber. Temporal characteristics of different cryosphere-related slope movements in high mountains. In *Proceedings of the 2nd World Landslide Forum*, volume 9, pages 285–297, Rome, Italy, Oct 2011. Springer.

