

Diss. ETH No. XXXXX

Design and Instrumentation of Environment-Powered Systems

A thesis submitted to attain the degree of

Doctor of Sciences of ETH Zurich
(Dr. sc. ETH Zurich)

presented by
LUKAS SIGRIST
M.Sc. ETH Zurich

born on 05.02.1992
citizen of
Affoltern am Albis ZH

accepted on the recommendation of
Prof. Dr. Lothar Thiele, examiner
Prof. Dr. Luca Mottola, co-examiner
Dr. Jan Beutel, co-examiner

2020



Institut für Technische Informatik und Kommunikationsnetze
Computer Engineering and Networks Laboratory

TIK-SCHRIFTENREIHE NR. XXX

LUKAS SIGRIST

Design and Instrumentation of Environment-Powered 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. XXXXX

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

Examination date: February 28, 2020

Abstract

Energy Harvesting presents a key technology to sustainably supply the billions of devices in the emerging Internet of Things (IoT). Converting physical signals such as radiation, temperature, vibration, etc. into electrical energy promises virtually unlimited energy to supply cyber-physical systems (CPSs) in a long-term and scalable manner. However, with an energy supply depending on a spatially and temporally variable environment significant non-determinism is introduced into the system.

In this thesis we explore the potential and limitations of supplying cyber-physical systems (CPSs) from environmental energy using only minimal energy buffering. We introduce novel design methodologies to supply applications reliably and efficiently, explore the energy yield of thermoelectric harvesting, and optimize the utility of data transmissions in infrastructure-less monitoring. Furthermore, we introduce a testbed and measurement support to assist designers in design aspects arising in energy harvesting systems.

Specifically, we make the following contributions:

- We introduce a novel measurement tool that combines high accuracy and portability. Enabling joint in-situ observations of the ambient, multiple energy flows, and application states, it provides critical insights during the design and verification of energy harvesting systems.
- We present a testbed for the emulation of radiation and temperature environments. In combination with a programmable, time- and event-triggered current sink, it enables fast and repeatable exploration, dimensioning

and validation of energy harvesting system design aspects.

- We introduce the first model for thermoelectric energy harvesting at the ground-to-air boundary that incorporates all components from the physical signal to the application. In combination with a newly proposed rectifier circuit, an optimized harvesting system is implemented. Extensive real-world evaluation attests the accuracy of the model and demonstrates unprecedented output power in the given harvesting scenario.
- We propose a novel energy management principle that decouples the energy harvesting and electrical load using a minimal energy buffer to allow each end to operate at its optimal operating point. An energy management unit (EMU) implementing this principle is designed and extensively evaluated. Efficient and reliable operation is demonstrated, even when the input power is significantly lower than the application requirements and exhibiting high variability.
- We study the utility of data transmitted in an infrastructure-less communication scenario supplied by energy harvesting. Using a model-based optimization approach, we derive a new data transmission scheme for long-term batteryless monitoring applications. Evaluation using a batteryless sensor nodes demonstrates accurate abstraction of the scenario using our model and significant gain in utility at minimal run-time overhead.

The methods and solutions presented are implemented and extensively evaluated under lab and real-world conditions. From these, we conclude that the methods and design tools presented enable efficient design and thorough evaluation of energy harvesting systems.

Contents

Abstract	i
List of Figures	ix
List of Tables	xvii
Acronyms	xix
1 Introduction	1
1.1 Challenges of Energy Harvesting	3
1.2 State of the Art	4
1.3 Thesis Contributions and Outline	7
2 Instrumentation of Energy Harvesting Systems	13
2.1 Introduction	14
2.2 Energy Harvesting System Design Flow and Related Work	17
2.2.1 Existing Measurement Devices	18
2.3 Measurement System Architecture	20
2.3.1 Architecture Overview	20
2.3.2 Precision Power Measurement	22
2.3.3 Mixed-Signal Capabilities	25
2.3.4 Measurement Control and Services	25
2.4 Performance Characterization	26
2.4.1 Characterization Setup and Calibration	27
2.4.2 Current Measurement Performance	28
2.4.3 Voltage Measurement Performance	30

2.4.4	Performance Summary and Positioning . . .	31
2.5	Case Studies	32
2.5.1	Low Power Optimization	34
2.5.2	Multi-Source Wearable Harvesting	35
2.5.3	Long-Term In-Situ System Evaluation . . .	37
2.6	Summary	39
3	Environment and Application Emulation Testbed	41
3.1	Introduction	42
3.2	Related Work	44
3.3	Physical Environment Emulation	46
3.3.1	Environment Emulation Requirements . .	47
3.3.2	Thermal Gradient Emulation	47
3.3.3	Visible Light Emulation	49
3.3.4	Solar Testbed Range and Accuracy	53
3.4	Electrical Load Emulation	55
3.4.1	Electrical Load Emulation Requirements .	55
3.4.2	SmartLoad Architecture	55
3.4.3	Electrical Load Evaluation	57
3.5	Automated System Evaluation	60
3.5.1	Automated Emulation Control	61
3.5.2	Controlling the Time Domain	61
3.5.3	Rapid Setpoint Characterization	62
3.6	Case Studies	63
3.6.1	Automated Transducer Characterization .	63
3.6.2	Energy-Triggered Application Execution .	66
3.7	Summary	67
4	Thermoelectric Harvesting from Natural Gradients	69
4.1	Introduction	70
4.2	Related Work	74
4.3	Thermal Energy Harvesting System	77
4.3.1	Thermal Harvesting Scenario	78
4.3.2	Thermal Harvester Architecture	78
4.4	Thermal System Model	79
4.4.1	Thermoelectric Generator Characteristics	81
4.4.2	Thermal Model of the Harvesting System	82
4.4.3	Radiation Model	83

4.5	Rectification and Electrical System Model	89
4.5.1	Low-Power/Low-Voltage Rectification . .	89
4.5.2	Electrical Load Matching and Voltage Conversion	91
4.5.3	Wireless Sensor Node Application	93
4.6	Environment Monitoring Use Case	94
4.6.1	System Dimensioning and Integration . .	95
4.6.2	Real-World Harvesting Performance . . .	96
4.7	Summary	104
5	Batteryless Energy Management	107
5.1	Introduction	108
5.2	Related Work	114
5.3	Energy Burst-Driven Task Execution	116
5.3.1	Predictable Executing using Energy Bursts	117
5.3.2	Efficient Task Based Execution	118
5.4	Energy Management Unit	120
5.4.1	Modeling Energy Flow, Buffer and Losses	120
5.4.2	Reducing Cold-Start Energy and Time . .	122
5.4.3	Minimizing Load Energy	123
5.4.4	Energy Management Architecture	124
5.5	Transient System Architecture	126
5.5.1	Low-Power Imaging Application (Load) .	127
5.5.2	Energy Burst Configuration	128
5.5.3	Software Execution Flow	129
5.5.4	Feedback Control for Energy Burst Scaling	129
5.6	Evaluation	131
5.6.1	Experimental Setup	131
5.6.2	Cold-Start Energy and Time Overheads .	133
5.6.3	Constant Input Power	136
5.6.4	Variable Input Power	137
5.6.5	Result Discussion	138
5.7	Summary	140
6	Long-Term Infrastructure-Less Monitoring	143
6.1	Introduction	144
6.2	Related Work	148
6.3	System Overview	151

6.3.1	Communication Scenario	151
6.3.2	Notation and Communication Model . . .	154
6.3.3	Sensor Node Architecture	157
6.3.4	Determination of the Policy Pool	158
6.4	Optimal Communication Policy	159
6.4.1	Data Utility	159
6.4.2	Probability of Data Reception	159
6.4.3	Optimization of Data Reception	162
6.4.4	Optimal Data Selection Policy	163
6.5	Model Generalizations	165
6.5.1	Relaxing Activation Interval	165
6.5.2	Aggregated Sensor Values	166
6.5.3	Energy Harvesting Awareness	167
6.5.4	Lossy Communication Channel	168
6.6	Sensor Node Implementation	168
6.6.1	Transient Sensor Node	169
6.6.2	Smartphone Receivers	172
6.7	Experimental Evaluation	172
6.7.1	Evaluation Setups	172
6.7.2	Sensor Node Characterization	175
6.7.3	Model-Based Scalability Analysis	176
6.7.4	Model Accuracy	177
6.7.5	Data Utility Analysis	179
6.7.6	Real-World Harvesting Use-Case	182
6.8	Summary	184
Appendices		185
6.A	Optimization Formulation Convexity	185
6.A.1	Convexity of Optimization Formulation . .	185
6.A.2	Convexity of Aggregates Extension	187
7	Conclusions	189
7.1	Contributions	189
7.2	Future Directions	192
Bibliography		195
List of Publications		221

Curriculum Vitae	225
-------------------------	------------

List of Figures

1.1	The interaction of the cyber and the physical worlds in an energy harvesting system, with new dependencies shown dashed. Compared to a traditional battery supply, the energy supply reaches from the physical into the cyber world and becomes dependent on the environment, resulting in a new, system-wide feedback loop in the form of an energy feedback cycle.	2
1.2	The high-level architecture of the energy supply in an energy harvesting cyber-physical system (CPS). The energy management links the variable input from the environment, an intermediate energy buffer, and satisfies dynamically changing application requirements.	4
1.3	The primary focus of each chapter of the design support and methodologies introduced in this thesis on the architectural landscape of an energy harvesting system.	9
2.1	Sample design aspects of harvesting-based systems. Boxes indicate aspects with distinct measurement requirements.	18
2.2	Overview of the RocketLogger system architecture. It consists of an analog front-end for high dynamic range power measurements, a BeagleBone hosting the management software and remote control, and external sensors for environment monitoring. . . .	21

2.3	High-level schematic of the RocketLogger’s current measurement circuitry. The feedback ammeter circuit is selectively switched on for high accuracy measurement of low currents.	23
2.4	Analysis of the transient burden voltage for different current steps from 0 mA to 500 mA, performed at time $t = 0 \mu\text{s}$. The results demonstrates a fast switching within $< 1.4 \mu\text{s}$, and a maximum burden voltage during that transition of 430 mV.	29
2.5	Comparison of the power traces for four alternative timed sleep implementations, sampled at 1 kSPS. The close-up on the right side shows the different sleep power levels and the varying wake-up delay of the three lowest power options.	34
2.6	The harvesting power levels and input voltages of a multi-source harvesting circuit combining solar and thermoelectric energy harvesting. The lower plot shows the ambient temperature and illuminance traces describing the time-varying harvesting scenario.	36
2.7	Power and ambient trace of a long-term in-situ validation of an energy harvesting-driven sensing system. Only this long-term in-situ test revealed the timing fault at $t = 5.97 \text{ h}$, since it is tied to the time-varying operating points of both, the source and load.	38
3.1	The architecture of the thermal testbed with two software controlled Peltier elements for consistent temperature gradient emulation. The adjustable distance enables evaluation of standalone thermoelectric generators (TEGs) or complete thermal systems including thermal guides.	47
3.2	The thermal testbed implementation. In this example, a thermal harvesting system consisting of a thermoelectric generator (TEG) and an attached copper thermal guide is evaluated in between the two Peltier elements.	48

3.3	The solar testbed consists of a dimmable LED light source, a mirror box for a uniform illuminance distribution and a Peltier element to control the solar panel's temperature. Temperature and illuminance sensors enable logging of the emulated conditions.	49
3.4	The fully integrated solar testbed. The top half represents the mirror-shielded light box with the light-emitting diodes (LEDs) mounted below the heat sink. The bottom of the light box consists of a Peltier controlled aluminum plate to maintain a constant temperature of the solar panel during the evaluations (hidden in the shadows). On the right side of the testbed the control circuitry is mounted: the LED dimmer with high dynamic range on top, the SmartLoad (presented in Section 3.4) on the bottom left, and the RocketLogger providing the measurement capabilities and hosting the testbed control (discussed in Section 3.5).	51
3.5	To provide control over a large illuminance range of 10 lx to 120 000 lx, a custom two-stage LED dimmer circuit was developed. It combines step-down current regulation and pulse width modulation (PWM) dimming. Multiple of these channels are provided for controlling a total of 36 power LEDs.	52
3.6	The illuminance emulation shows a high accuracy of $8.8\text{ lx} + 0.5\%$ over the range from 10 lx to 120 klx. A significant increase in emulation error below 100 lx shows the limitations of the presented testbed for extremely low illuminance levels. . . .	54
3.7	High-level schematic of the SmartLoad programmable current sink. A feedback controlled MOSFET sinks a current that is set by the digital-to-analog converter (DAC). The controller exposes an I ² C interface for configuration and replays current traces in time- and event-triggered manner.	56
3.8	The input voltage step rejection evaluation of the SmartLoad shows very short current spikes of about $0.3\ \mu\text{A}$ at the lowest current range.	59

3.9	Combining the different emulation tools with mixed-signal measurements and a central control enables a streamlined system emulation flow. . . .	60
3.10	The setpoint characterization measurement of an AM-5412 solar cell shows drastically increased error if the analysis is performed using data points too close to the setpoint change.	62
3.11	The automated transducer characterization enables rapid and exhaustive evaluation over a wide range of operating points. The illustrated characterization of an AM-5412 photovoltaic cell at a temperature of 25 °C consists of 2 160 setpoints that were recorded in only 76.6 min.	64
3.12	The thermal capacitances of the testbed and limited output power of the testbed demand for 30 h for a 461 setpoint characterization of a thermoelectric generator (TEG). The steady state detection guarantees accurate measurements without transient effects originating from setpoint switching. . . .	65
3.13	The experimental setup used for the energy management unit (EMU) evaluation. The SmartLoad replays the current trace previously recorded using the original application circuit.	66
3.14	Comparison of the power traces of the original application and its emulation with the SmartLoad synchronized at the application trigger time. . . .	67
4.1	The rock temperature profiles from an alpine rock wall scenario illustrate the delayed and attenuated propagation of thermal energy through the rock material.	79

4.2	The architecture of the thermoelectric energy harvesting platform for the ground-to-air boundary: from the radiation regime on the ambient side (a), to the thermal guide connecting the ground at depth (b), using a TEG transducing the heat flux into electrical energy (c), and the electrical subsystem for rectification (d) and up-conversion of the low voltages to charge a battery (e) and supply the wireless sensor node application (f). . .	80
4.3	The components of the lumped parameter model and their corresponding mechanical parts in the thermal guide design. The figure shows the dimensions resulting from the design choices made for the case study detailed in Section 4.6.	84
4.4	Comparison of testbed experiments and simulations for the thermal guide model, including the characterized TEG properties. The overshoot of the forced temperature is attributed to the step response of the control circuit in Peltier devices. .	85
4.5	A representative two days excerpt of the model verification experiment in an urban concrete wall deployment. The measurements show the observed model inputs for the <i>thermal guide only</i> and the <i>end-to-end model</i> on the top. The comparisons of the models in terms of the TEG gradient ΔT_{TEG} in the middle and the matched output voltage V_{TEG} in the lower plot attest a high accuracy of the end-to-end model and its thermal guide component. . . .	86
4.6	The low-voltage rectification circuit is based on a bridge of four enhancement mode nMOSFETs and a low-power control circuit.	90
4.7	The short-term SPICE simulation confirms the functional correctness of the proposed MOSFET-based low-power low-voltage rectification.	91

4.8	Left: exploded view of the harvesting system components consisting of the thermal harvester, TEG, rectification circuitry, and the wireless sensor node. Right: the system deployed in the rock wall for the long-term case study (photograph of the SE position 29).	97
4.9	Long- and short-term (3 day differences) battery voltage evolution showing steady charging at the south locations.	99
4.10	Ten days of temperature difference and the resulting harvesting power showing the influence of weather and location.	100
4.11	Histogram of the temperature gradients and harvested energy during the 237 days case-study of the SE facing sensor node deployment at position 29. .	101
4.12	Histogram of the temperature gradients and harvested energy during the 625 days case-study of the NW facing sensor node deployment.	102
5.1	The proposed components (dashed) for efficient and reliable transient system design: an energy management unit (EMU) exposing a control interface for feedback-based dynamic energy burst scaling (DEBS).	112
5.2	Illustration of the energy management unit (EMU) operation principle: accumulating the variable input energy until the required energy can be provided in an <i>energy burst</i> to the load for task execution. Starting from a fully depleted buffer, demands first reaching the load's minimal supply voltage, leading to a cold-start energy and time overhead (time axis not to scale).	117
5.3	Splitting the application into tasks τ_1 and τ_2 allows separate execution of them at their respective optimal power point using dynamic energy burst scaling (DEBS). This reduces the energy footprint and increases the responsiveness of the application.	119

5.4	The boost/buck architecture of the energy management unit (EMU) with a control circuit exposing the energy burst configuration to the load. An optimally sized capacitor serves as energy buffer for reliable task execution and at the same time reduces the cold-start overhead.	124
5.5	The proposed architecture integrates the energy management unit (EMU) with an application circuit consisting of an MSP430FR5969 microcontroller and a Centeye Stonyman image sensor. In the evaluation it is powered from a solar panel.	127
5.6	The software execution flow of an application making use of the dynamic energy burst scheme.	129
5.7	Cold-start time and energy analysis for different input power levels.	133
5.8	Evaluation of the system efficiency η_{sys} under constant input power conditions.	134
5.9	Evaluation of the application execution rate U_{exp} under constant input power conditions.	135
5.10	Time domain comparison of the real-world experiment using dynamic energy bursts and the corresponding simulation taking the observed harvesting power P_{in} as input.	139
6.1	The considered scenario consists of a battery-less sensor node monitoring the environment and broadcasting (RF) data packets of N sensor values to smartphones with listening intervals described by distribution $\tilde{\mathbf{R}}$. The application specification $\tilde{\mathbf{Q}}^{\text{spec}}$ describes the relative importance of individual data units.	151
6.2	Examples of the relative importance specification $\tilde{\mathbf{Q}}^{\text{spec}}$ for a few reference scenarios with differing requirements.	153
6.3	Examples of receiver characteristics $\tilde{\mathbf{R}}$ for three different scenarios.	153

6.4	The overall architecture consists of an <i>online</i> harvesting aware data transmission scheme. The key component is a pool of communication policies that are optimized <i>offline</i> based on a model for the infrastructure-less communication scenario.	155
6.5	Enumeration of the scenarios where i time unit old data is not received (s_0, \dots, s_i)	161
6.6	The transient Bluetooth Low Energy (BLE) sensor node is based on the CC2650 system on a chip (SoC). The SoC interfaces a real-time clock (RTC) with separate backup power domain and a ferroelectric RAM (FRAM) to provide timekeeping capabilities and data retention across power failures. For batteryless operation, the node relies on the energy management unit (EMU) and energy burst principle introduced in Chapter 5 and a solar panel as energy source.	169
6.7	The custom designed transient sensor node used to implement and evaluate the communication scheme has a small size of only 35 mm \times 53 mm, about the size of the match-box show in the background. It is powered from the solar panel mounted on the back of the node.	170
6.8	Problem specification with utility specification \tilde{Q}^{spec} and a stochastic listening interval distribution \tilde{R}	173
6.9	The model accuracy experiment results with $\delta = 1$ demonstrate a high model accuracy.	178
6.10	The data utility provided by policies optimized for specific values of λ and the baseline scheme in comparison to the optimally achievable data utility.	179
6.11	Comparison of the data utility of the baseline and optimal scheme for a problem specification with a history coverage shortened to 25 s.	180
6.12	Real-world harvesting use-case results with $\delta = 1$: high model accuracy despite additional sources of error.	182

List of Tables

2.1	Performance characteristics of the RocketLogger. Noise and bandwidth values are shown for 1 kSPS and (64 kSPS).	31
2.2	The RocketLogger is positioned between stationary digital power meters and embedded power loggers.	33
3.1	DC performance of the SmartLoad.	58
3.2	Dynamic performance of the SmartLoad.	59
4.1	Performance Comparison of Different Switching Approaches for Low-Power Low-Voltage Rectification.	92
5.1	Energy burst requirements of the considered imaging application when executing its tasks separately, and when executing the entire application in a single burst.	128
5.2	System performance analysis for both execution profiles in the real-world experiment with variable input power.	137
6.1	Transient Bluetooth Low Energy (BLE) sensor node operation characteristics in terms of energy and time: mean of 12 819 activations, including the 95 % percentile of the absolute deviation from the mean.	176

6.2	The model error analysis for the model accuracy experiment shows very small errors for all power levels pwr^{in}	177
6.3	The model errors analysis for real-world harvesting use-case experiment in an office deployment with illuminance levels varying from 300 lx to 800 lx demonstrates high model accuracy. The statistically not significant results for the $\delta = 2$ were omitted.	183

Acronyms

ADC analog-to-digital converter.

BLE Bluetooth Low Energy.

CPS cyber-physical system.

DAC digital-to-analog converter.

DEBS dynamic energy burst scaling.

DVFS dynamic voltage and frequency scaling.

EMU energy management unit.

FRAM ferroelectric RAM.

IoT Internet of Things.

I²C Inter-Integrated Circuit.

LED light-emitting diode.

MOSFET metal-oxide-semiconductor field-effect-transistor.

MPPT maximum power point tracking.

op-amp operational amplifier.

PCB printed circuit board.

POR power-on-reset.

PRU programmable real-time unit.

PWM pulse width modulation.

RAM random-access memory.

RF radio frequency.

RFID radio-frequency identification.

RMS root mean square.

RMSE root mean square error.

RTC real-time clock.

SMU source measure unit.

SoC system on a chip.

SPDT single pole, double throw.

SPICE Simulation Program with Integrated Circuit Emphasis.

TEG thermoelectric generator.

USB Universal Serial Bus.

WSN wireless sensor network.

1

Introduction

The advent of microelectronics has enabled processing of information at massive scale. As part of CPSs that interact with the environment through sensing and actuation, information processing builds the foundation of a vast number of ubiquitous applications [RLSS10]. Leveraging wireless communication, sensing systems are nowadays deployed in the form of wireless sensor networks (WSNs) in many scenarios as they require only minimal infrastructure during operation and are therefore installed with minimal effort. In the emerging Internet of Things (IoT) the instrumentation of our environment and everyday life is not only expected to continue but to grow faster than before [Sta14].

The majority of these systems rely on a primary battery as energy supply. Despite the considerable advances in low-power system design, this type of storage will run out of energy eventually and needs to be replaced. Energy harvesting is seen as a key technology to provide a long-term and sustainable energy supply that scales to countless number of




Figure 1.1: The interaction of the cyber and the physical worlds in an energy harvesting system, with new dependencies shown dashed. Compared to a traditional battery supply, the energy supply reaches from the physical into the cyber world and becomes dependent on the environment, resulting in a new, system-wide feedback loop in the form of an energy feedback cycle.

devices [PS05]. Extracting energy from physical signals such as light, temperature, movement, vibration, electromagnetic waves, etc., it supplies applications directly from the environment in which it is deployed [BASM16].

This paradigm shift in energy supply is also reflected in the interactions of the cyber-physical feedback cycle depicted in Figure 1.1. With harvesting, the supply of energy becomes strongly dependent on the environment and demands continuous adaptation to the variable environmental conditions. As a result, the previously purely physical energy supply from primary batteries is extended or even replaced by a complex hardware-software system that presents an additional link between the physical and cyber worlds. Ultimately, the energy flow establishes its own feedback cycle, emphasizing that it becomes a first order concern in the design of efficient energy harvesting systems.

1.1 Challenges of Energy Harvesting

Related to the above discussed dependencies and feedback loops, energy harvesting adds distinct challenges related to system design that need to be addressed to build functional and efficient systems.

Non-Deterministic Environment. The energy extracted by transducers such as solar panels or thermoelectric generators (TEGs) is coupled to the physical signals of the environment. These signals exhibit significant and typically non-deterministic spatial and temporal variability. Through energy harvesting, this non-determinism has a direct influence on the system's operation. Consequently, any energy harvesting system has to control and/or mitigate this variability.

Highly Variable Load Behavior. Similarly, today's cyber-physical systems (CPSs) exhibit a high variability in power consumption. Due to adaptive techniques used in low-power design, such as dynamic voltage and frequency scaling (DVFS), numerous system operation and sleep modes, power gating, and aggressive duty-cycling, they can operate in many different configurations with different functional properties and energy requirements [BdM00]. While the operating point might be controlled by the system designer, selecting the optimal set of often dynamically adapted operating points presents a non-trivial task in itself. The actual load behavior might even depend on the non-deterministic physical signals it observes, for example when employing an event-triggered sensing methodology for energy efficient monitoring.

To maximize system performance, the energy input has




Figure 1.2: The high-level architecture of the energy supply in an energy harvesting cyber-physical system (CPS). The energy management links the variable input from the environment, an intermediate energy buffer, and satisfies dynamically changing application requirements.

to be maximized and the load's energy required per unit of work minimized. In energy harvesting systems, maximization of the energy input demands for continuous adaptation to the variable environment conditions. The load uses dynamic power management to minimize its energy consumption, which results in highly dynamic energy requirements due to different operating points. Addressing these incompatible requirements in a joint systems design represents a major challenge. Only novel energy management methods that enable decoupling of these concerns and using intermediate energy buffering can simultaneously address the requirements of both sides. These consequences are also reflected in the typical energy harvesting system architecture shown in Figure 1.2 that is conceptually found in any energy harvesting system.

1.2 State of the Art

The concept of energy harvesting has been around for many years. In the following we provide a broad overview on the

state of the art in supplying cyber-physical systems (CPSs) from small-scale, in-situ energy harvesting.

Energy Harvesting Methods

As mentioned earlier, energy can be extracted from many physical phenomena in the environment. Commonly exploited sources include light using photovoltaic cells [RKH⁺05], temperature gradients with TEGs [Sny09], vibrations using piezoelectrics [KPK12, MYR⁺08] or radio frequency (RF) signals [SCH16]. An exhaustive survey including additional energy harvesting sources is given in [BASM16]. To exploit the diversity of sources for a more continual energy input, methods for combined harvesting from multiple sources have been explored on circuit and system level [BC12, WMM⁺13]. With the energy stemming from various physical signals, each source has its distinct characteristics and requirements to be addressed for efficient energy extraction.

Efficient Energy Extraction

To maximize energy transfer, the electrical (output) impedance of transducers needs to be matched. For transducers that exhibit an input dependent impedance such as solar cells, maximum power point tracking (MPPT) is typically employed to closely track the optimal electrical operating point of transducers [FPSV05, BBMT08, KLYK09]. For many transducers also their physical interface requires to be matched appropriately to the environment, for example the thermal system to maximize the temperature gradient for thermoelectric harvesting [Sny09] or the mechanical system to match the resonance frequency of piezoelectric elements [KER⁺10].

While previous work typically focuses on specific aspects

of energy harvesting, they have to be optimized jointly across the cyber-physical boundary to maximize the extracted energy.

Energy-Neutral Systems

Ideally, energy harvesting can provide the energy required for sustained system operation. Energy neutrality requires that the average harvested energy is sufficient to satisfy application requirements. However, also the variability of the energy input needs to be mitigated. The approaches taken to enable energy-neutral systems are power subsystem planning [KHZS07, BSBT14b], and run-time adaptation of the application execution [VGB07, BSBT14a, ABD⁺19]. However, if the input power is very low or the energy buffer constrained, energy-neutrality may not be guaranteed, resulting in intermittent system operation.

Intermittent Operation

When miniaturizing the systems, motivated by form-factor and cost considerations, the energy input is likely too little and energy buffers are too constrained for energy-neutral operation. To address the resulting intermittency of the energy supply, several mechanisms for saving system state across supply failures have been proposed. For example, non-volatile processors avoid volatile system state by design [MLS⁺16]. Various mechanisms to back up volatile state before supply failure have been described on the hardware [STSG16, HFdB17] and software level [RSF11, BM17], or a combination of both [JRLR15, BWD⁺16]. Problems related to the data consistency in checkpoints were addressed by data management schemes such as [CL16, Hic17] and programming models like [MCL17, HSS17]. To efficiently adapt to the varying application requirements, dynamic capacity

adaptation [CRL18] and distributed energy buffering [HSS15] have been proposed.

These approaches primarily focus on tolerating the frequent energy supply failures caused by the variable environment. However, the broader challenge of reliable and efficient operation under these constrained energy input conditions demands addressing the environment's and load's variability jointly.

Design Support for Energy Harvesting

To support developers in the design of energy harvesting system, several tools have been introduced. They provide emulation of the electrical characteristics of transducers [HSSS17], reproduction of harvesting circuits behavior in distributed systems [GCZ19], and the forcing of physical signals, such as light, to reproduce consistent harvesting system behavior [HRB⁺17]. For thorough evaluation of energy harvesting CPS the main sources of non-determinism, namely the physical environment and application requirements, should be consistently reproduced with exhaustive coverage of their full variability.

1.3 Thesis Contributions and Outline

In this thesis we explore the potential and limitations of supplying cyber-physical systems from energy harvesting when using only minimal energy buffering. We develop design methodologies to reliably execute applications with variable requirements under constrained and non-deterministic energy input from the environment, and we study the implications of a tiny energy buffers. In parallel, we present novel tools that support specific requirements in the design space

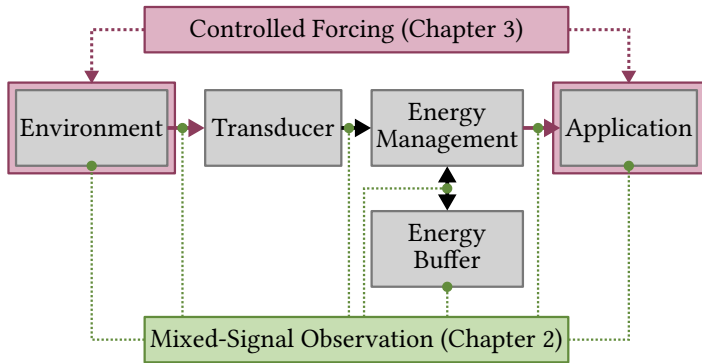
exploration, system dimensioning, evaluation and verification of energy harvesting systems. Figure 1.3 highlights the focus of each chapter within the architecture of energy harvesting systems, separated by the contributions in terms of design methodologies and tools. In the following we present the outline of this thesis with an overview of the contributions in each chapter.

Instrumentation of Energy Harvesting Systems (Chapter 2)

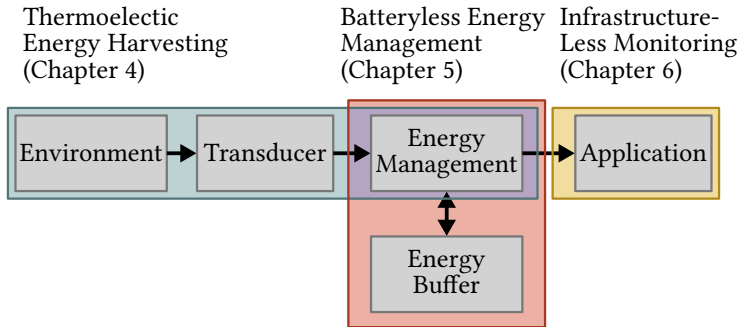
We introduce the RocketLogger, a novel measurement device providing joint observation of the environment's physical signals, of multiple energy flows, and of the digital state. The mixed-signal capabilities significantly extend the observability of the many aspects in the energy feedback cycle of energy harvesting cyber-physical systems. The detailed observation capabilities are also provided in-situ due to the portability of the feature-rich measurement device.

Testbed for Environment and Application Emulation (Chapter 3)

We present a testbed that reproduces thermal and radiation properties of the environment, and is capable of sinking arbitrary current profiles in time- and event-triggered manner. The controlled forcing of specific properties in the energy flow enables detaching specific design aspects from the energy feedback loop in cyber-physical systems. This offers evaluation, exploration and dimensioning of energy harvesting systems and their components in isolation, while still using an emulated but realistic physical environment. Providing flexible and well-controlled scaling of the time domain, the testbed enables fast and reliable experimental feedback in the design of energy harvesting systems.



(a) Design support contributions.



(b) Focus of the methodological contributions.

Figure 1.3: The primary focus of each chapter of the design support and methodologies introduced in this thesis on the architectural landscape of an energy harvesting system.

Thermoelectric Energy Harvesting from Natural Temperature Gradients (Chapter 4)

We introduce the first model for thermoelectric energy harvesting at the ground-to-air that incorporates all aspects from the environment to the application, and extends significantly to the physical aspects of the specific energy source. In combination with a newly proposed rectifier circuit, an optimized energy harvesting system is implemented and extensively evaluated in real-world deployments. The evaluation attests the accuracy of the model and demonstrates unprecedented output power in comparable harvesting scenarios.

Efficient Energy Management for Batteryless Harvesting Systems (Chapter 5)

We propose a novel energy management principle that decouples the efficient extraction of energy from the non-deterministic environment and the energy supply requirements of the load. This enables both ends to operate at the respective optimal operating point, and hence maximizes the application performance. At the same time the method allows mitigation of the energy input variability stemming from the environment using a minimized energy buffer. An energy management unit (EMU) implementing this principle is designed and evaluated extensively. The results demonstrate the feasibility of supplying applications reliably and with high efficiency, even under very low or highly variable energy input.

Infrastructure-Less Monitoring with Batteryless Sensors (Chapter 6)

We study the utility of data transmitted in an infrastructure-less communication scenario. Following a model-based opti-

mization approach, we derive a new data transmission scheme for long-term batteryless monitoring applications. The implementation and evaluation of the scheme on batteryless sensor nodes demonstrate that the model provides accurate abstraction. Further, it shows that the proposed scheme results in significant gain in data utility at a very low run-time overhead.

2

Instrumentation of Energy Harvesting Systems

As energy harvesting becomes ever more prevalent in embedded systems, designers require appropriate tools to be able to test, validate, and characterize the performance of their systems in spatially and temporally variable environments. The need for real-world experiments demands for a compact and portable solution to perform accurate in-situ measurements and environmental logging. Responding to this demand, we introduce the RocketLogger, a portable measurement device capable of accurately measuring power at source and sink, exhibiting a high dynamic range and logging environment conditions in long-term in-situ deployments. To address the challenge of measuring the large dynamic range typical for energy harvesting and low-power systems with minimal impact, we propose a combination of a shunt resistor and feedback ammeter circuit with seamless switching. We extensively characterize the measurement performance of the

RocketLogger and demonstrate the advantages of its trade-off between high-performance lab equipment and highly mobile power loggers in three use-cases.

2.1 Introduction

The advances in low power embedded systems over the past decades have considerably increased the lifetime of battery powered devices. However, deployments in hardly accessible areas and the billions of devices in the emerging Internet of Things (IoT) demand for long-term deployments with virtually unlimited lifetimes. Battery-only designs are not an option, since their limited lifetime requires expensive maintenance. Energy harvesting and ultra-low power system design are seen as key solutions to that problem and gained increasing research attention in recent years.

Energy Harvesting Design Aspects

Unfortunately, exploiting energy harvesting adds another layer of complexity to the design process. Variable and application-specific environmental conditions have a direct impact on parameters like harvested energy and power conversion efficiencies. Harvesting-based systems need to handle a wide range of input power levels from nW to mW and adapt to changing supply conditions while efficiently using available energy. The experience with a broad range of systems like remote sensing [BBF⁺11], harvesting system design [MTBB10], dimensioning [BSBT14a], and management [GPB⁺15, ABD⁺19], and wearables [MBS⁺16, TSM⁺17] shows that they share a common design problem. Regardless of the design strategy used, a precise characterization of the environment-dependent energy budget, as well as simulta-

neous characterization and optimization of the application's energy consumption is inevitable.

The design process of these systems typically starts with roughly estimated environmental variables like luminosity, temperature differences, vibration energy or pressure changes. Once an initial prototype is ready for testing, different metrics such as harvesting efficiency, the storage element's capacity and leakage, and active/sleep currents of individual components can be measured to provide feedback to the design flow. These measurements, however, include widely ranging currents in the nA- μ A range for quiescent/sleep operation modes, and up to 100's mA for active modes. To characterize a harvester's behavior in its final placement in the environment, as opposed to the lab, measurement tools must be deployed with it. While measuring power and environmental properties on their own are well known problems, portable tools that accurately and reliably measure harvesting-based systems do not exist.

Measurement Requirements

Without harvesting, low power systems can be validated using real-time cycle-accurate energy measurements in the lab, such as in [CKL00]. Now system designers need to employ novel tools to design, test, and evaluate not only embedded systems themselves but also energy harvesters and the environment they operate in. There are many challenges involved in building these tools. Multiple voltage and current channels with high dynamic range are required to simultaneously measure energy harvesting and consumption. To link these measurements with the environment in which the system operates, the physical properties related to the energy harvesting method have to be monitored. Furthermore, support for tracing digital signals is needed to observe the state of

the application consuming the harvested energy. The sample rate of these signals has to be sufficiently high to trace the system operation in detail. If dynamic range switching is used, it needs to operate much faster than the data acquisition as intermediate samples will be inaccurate otherwise.

As the mixed-signal measurements are deployed the field, they need to be sustained for potentially long periods of time. At the same time, the measurement equipment must run independently from the load being measured, influence it as little as possible, and have large memories for storing long-term measurements and/or network connectivity. Otherwise, the device will either run out of memory quickly, not work in adverse power conditions or significantly degrade the system's harvesting or power conversion efficiency. Consequently, the measurement device must be portable, rechargeable, and minimize measurement leakage currents and burden voltages to reduce their impact on the measured device.

Contributions

In this chapter we introduce the RocketLogger, a hand-held measurement device with the diverse sensing capabilities designed especially for harvesting-based system design. Specifically, we address the problems of the extreme dynamic range of current by combining two separate current measurement circuit methods using fast and seamless switching among them. Thanks to the logger's portable design and remote configurability, it enables long-term in-situ measurements. It provides a total of four voltage and two current measurement channels with high accuracy and the ultra-high dynamic range necessary to characterize energy harvesting and application circuits. In addition, it incorporates tracing of six digital signals and measures the environmental conditions in which the system operates.

The RocketLogger is the first device to balance the trade-off between a full-featured, high-performance lab measurement equipment [Kei13a] and lower fidelity, mobile power measurement circuits [ZX13], thus satisfying essential needs of the system design process for energy harvesting-driven applications. The custom analog front-end allows logging at sample rates up to 64 kSPS. The fast range switching within 1.4 μ s guarantees continuous power measurements starting from 4 pW at 1 mV up to 2.75 W at 5.5 V. We extensively characterize the RocketLogger's performance and demonstrate its usefulness in three use-cases at different stages of the system design flow.

Outline

The remainder of this chapter is organized as follows: We cover measurement related aspects of the energy harvesting system design and existing solutions in Section 2.2. Section 2.3 introduces our proposed measurement architecture and the measurement circuit details. Subsequently, we characterize the devices measurement performance in Section 2.4. Three sample case studies of how the RocketLogger supports the design of energy harvesting devices are discussed in Section 2.6, before we summarize our contributions and findings in Section 2.6.

2.2 Energy Harvesting System Design Flow and Related Work

Through energy harvesting, sensor nodes can reach energy neutral operation. However, there are several design parameters which have a direct impact on the node's self-sustainability. For example, a transducer's harvesting and

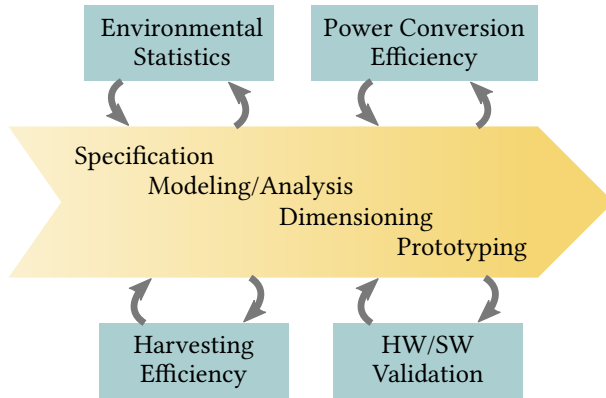


Figure 2.1: Sample design aspects of harvesting-based systems. Boxes indicate aspects with distinct measurement requirements.

power conversion efficiencies depend on environmental conditions such as light, movement, temperature or pressure changes. In many cases, such as wearable or implantable devices, the environment is not an easily-predicted macroscopic variable. Extensive in-situ measurements are thus required to gather statistically representative datasets, which can then be used in a typical embedded system design flow, as shown in Figure 2.1. Several phases of the design process require diverse measurements: from environmental statistics via harvesting/conversion efficiencies and dimensioning to validating the hardware and software of different components or prototypes. Currently available measurement equipment does not satisfy all the requirements of these different, but related design aspects.

2.2.1 Existing Measurement Devices

There is a wide range of power measurement devices with varying degrees of size and functionality. On one side there

are lab measurement devices that provide fast and highly accurate measurements, like the Keithley 2000 and Keysight 34400 precision digital multimeter series [Kei13a, Key19]. However, these devices are bulky, require standard high voltage AC supply and have a limited measurement buffer depth. More recently, some specific needs of the energy harvesting based community have been addressed with new tools. These tools can be loosely classified in to following categories.

In-Situ Energy Monitors

These are relatively simple measurement circuits, ranging from shunt resistors with differential analog-to-digital converter (ADC) [ZX13, JDCS07, RSW19] to coulomb counters [DFPC08, NPPS16]. While devices can have relatively wide measurement ranges, their main design goal is to bring basic energy awareness to sensor nodes through a single current channel. Since they are tied to the load, they cannot work independently or in adverse power conditions and they cannot provide long-term statistics.

Standalone Embedded Power Loggers

These devices are more complex than online energy monitors since they are independently powered observers with a wider dynamic range and the capability to store or forward data to a host for long-term logging. This was the approach taken by [AH09] and [PBLS14] to log the energy consumption of wireless embedded systems. In [BZL14], energy traces are used in conjunction with a kernel event logger to perform detailed power analysis in smartphones. These systems, however, are designed to measure battery-based nodes with voltage regulation.

Distributed Profiling Testbeds

These platforms are also external devices which observe different parameters pertaining to one or more node. Flocklab [LFZ⁺13], for example, can trace logical events, actuate GPIO pins, and profile the power consumption in a distributed manner. Similarly, SmartEye [PGN12] is a testbed designed for in-situ measurements of IoT nodes, with similar capabilities but with self-sustainability in mind. Very recently a testbed architecture to extend the recording and actuation to energy harvesting has been proposed in [GCZ19]. While we borrow the idea of mixed-signal capabilities for detail in-situ observation these platforms introduced, they target at distributed evaluation of wireless embedded systems.

2.3 Measurement System Architecture

As previously discussed, measuring and validating harvest-based systems is a non-trivial task. Their design incorporate many specifically optimized components that have to be taken into account from an early design stage. To instrument and validate these designs, portable and low-cost measurement equipment with multiple, highly accurate voltage and current channels covering a wide dynamic measurement range is required. Measurement devices that satisfy this feature combination, including seamless auto-ranging and minimal load impact, are not commercially available.

2.3.1 Architecture Overview

To address the lack of measurement equipment satisfying the specific requirements of energy harvesting system design, we introduce a novel mixed-signal measurement device architec-

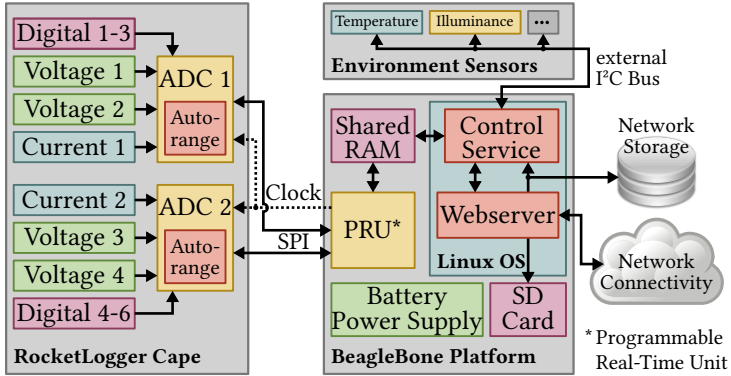


Figure 2.2: Overview of the RocketLogger system architecture. It consists of an analog front-end for high dynamic range power measurements, a BeagleBone hosting the management software and remote control, and external sensors for environment monitoring.

ture. We present a novel analog front-end for high accuracy power measurements with a very high dynamic range. Different digital sensors can measure environmental variables such as temperature, illuminance and atmospheric pressure, using a digital bus. The necessary control infrastructure is implemented on top of the Linux operating systems to handle the large volume of data and to manage the synchronous acquisition of the various measured variables. To maintain a hand-held form factor, the BeagleBone Green embedded Linux platform was selected [Bea]. Remote user interface and data storage features were implemented as Linux applications to enable remote control and observation in long-term in-situ measurements. An overview of the architecture and its components is shown in Figure 2.2. In the following, the requirements and challenges of the implementation is discussed for each architecture component.

2.3.2 Precision Power Measurement

Due to the large dynamic range of harvesting power and low power system consumption, a single circuit cannot cover the entire range. We address this issue by combining two current measurement circuits, one optimized for precise measurement of low current and the other for high current measurements. Integrating them into a combined current measurement channel demands a mechanism that switches seamlessly between them to not impacting the device under measurement.

Low Current Measurement

Measuring the current generated/consumed by harvesting/application circuits is essential for early design decisions and prototyping. Low power harvesting scenarios and the much improved sleep current of today's microcontrollers demand current measurement ranges down to the 10's nA. To measure these low currents accurately with minimal impact on the device under test, a feedback ammeter circuit [Kei13b, Ch. 1] is used. This operational amplifier (op-amp) based circuit allows measurement of ultra-low currents in the nA range because of its high amplification for the feedback resistance R_{fb} of 680 Ω . The voltage drop that can be measured with an analog-to-digital converter (ADC) is generated at the output of the op-amp. This has the advantage of virtually zero burden voltage at the input, therefore minimizing the disturbance on the device under test. However, the measured current flows through the op-amp and is therefore limited by its output current capabilities. To guarantee ultra-low current measurements down to 10's nA, the RocketLogger's low current measurement circuit is limited to a maximum range of ± 2 mA. Beyond this range the feedback ammeter needs to be bypassed, because it introduces an uncontrollably high burden voltage when the op-amp saturates. To handle

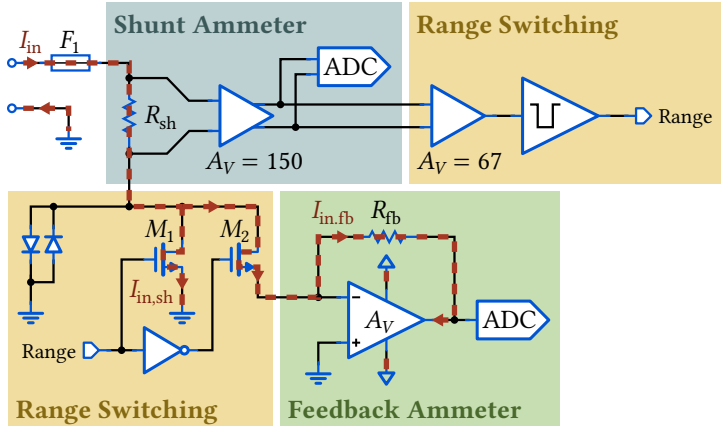


Figure 2.3: High-level schematic of the RocketLogger’s current measurement circuitry. The feedback ammeter circuit is selectively switched on for high accuracy measurement of low currents.

this, we introduce the range switching circuit discussed below.

High Current Measurement

In order to measure currents above 2 mA, a shunt ammeter circuit [Kei13b, Ch. 1] is employed. A small shunt resistance R_{sh} of 50 m Ω keeps the burden voltage low, thus minimizing the impact on the device under test. The voltage drop across the resistor is amplified with an instrumentation amplifier before analog-to-digital conversion. This allows accurate current measurement in the range of 2 mA to 500 mA with a measurement noise floor in the 10’s μ A range.

Seamless Measurement Range Switching

To achieve a total measurement range of 10 nA to 500 mA, the RocketLogger needs to seamlessly switch between the two circuits. To guarantee minimal disturbance of the measured

device during fast current transients, the range switching is integrated in the analog front-end as shown in Figure 2.3. This ensures that the feedback ammeter is deactivated quickly during large current steps, e.g. at waking up from sleep states, before it leads to a large burden voltage due to saturation. To detect the measurement range, an amplified version of the shunt ammeter's output signal is used: even if this signal is dominated by measurement noise for low currents, it remains valid and can be used to detect thresholds that are two orders of magnitude higher than the noise floor. The threshold detection is implemented using a window comparator. In low current range of ≤ 2 mA the MOSFET M_2 shown in Figure 2.3 is activated to pass the current also through the feedback ammeter. For higher currents the feedback ammeter is bypassed using M_1 to avoid a high burden voltage in case of saturation. The shunt ammeter is always active to provide a valid input for the range switching circuit. In addition to digitizing the current measurement circuit outputs using an ADC, the range switching signal driving the MOSFETs is recorded. This allows merging the two current measurement ranges in the analysis.

Voltage Measurement

While logging the voltage is not mandatory for battery powered devices with regulated voltage supply, this becomes a necessity for energy harvesting systems. Both, the harvesting and application circuit, can have time-varying operating points that directly impact the system's efficiency. Compared to the current measurement described above, voltages do not cover an extremely large measurement range. Voltages of 10's μ V up to a few volts can be measured directly using available precision ADCs. However, isolation from the parallel current measurements is a necessity to not disturb the current

measurements. This is achieved using voltage buffers with ultra-low input leakage currents and appropriate shielding of the measurement probes.

2.3.3 Mixed-Signal Capabilities

In parallel to current and voltage measurements, the Rocket-Logger supports logging of up to 6 digital inputs for tracing the state of the system under test. These measurements are complemented by environmental logging using temperature, luminosity, pressure, or vibration sensors that can be connected to a digital sensor bus. Additional environment sensors can be integrated with little effort by connecting them to this external accessible bus. Logging the system state and environment is essential, not only for characterization of the harvesting sources, but also for tracing the execution state of a sensor node.

2.3.4 Measurement Control and Services

Control and observation capabilities of the measurement device present an important, particularly when such devices are deployed in-situ for long-term evaluation.

Acquisition Control

The measurement control and coordination is performed in a background service on top of the Linux operating system. While lower rate, digital environmental sensors are read out directly by this control service, the power measurement data is not. Sampling the analog and digital inputs at up to 64 kSPS to track task executions in the sub-millisecond range results in up to 13 Mbit/s of measurement data. Because the read-out of these ADC conversion results needs to be

done reactively upon data availability with very low latency, the programmable real-time unit (PRU) of the BeagleBone core is used for timely data transfer. This data is buffered in the shared random-access memory (RAM), from where it is further processed by the control service. The service then combines the electrical with the environment measurements and applies calibration, before storing the measurements in the file system or forwarding them to an interactive display.

High Level Services

A data logger for harvesting-driven system design is required to be portable for in-situ measurements. This in turn demands for remote control and management of data acquisition. Building on top of wired or wireless connectivity provided by the underlying operating system, the RocketLogger exposes a web interface for remote control. It allows configuring measurement settings like sample rate and channels to record, as well as starting and stopping the acquisition and downloading completed measurements. In addition, the web interface also provides online preview of running measurements. This allows checking for correct measurement setup easily and tracking of the measurement progress and system state during long-term data acquisitions. The network also allows offloading measurement data to a network storage. This is especially important for long-term measurements, as the local storage on the BeagleBone may eventually be full.

2.4 Performance Characterization

This section presents the performance characterization of the RocketLogger measurement solution. Following the introduction of the experimental setup used for these measurements

and the discussion of the individual performance metrics, we summarize the RocketLogger's feature set and its unique positioning with regards to existing measurement devices.

2.4.1 Characterization Setup and Calibration

Prior to performance characterization, the RocketLogger was calibrated. For the linearly designed channels, a first order calibration consisting of offset and gain is used. A Keithley 2450 source measure unit (SMU) was configured to supply stepwise increasing current and voltage inputs to the RocketLogger. Steps of 2 mA for high current, 20 μ A for low current range, and 100 mV for the voltage channel were used for the calibration. The calibration measurements were then averaged for each current and voltage channel, and the channel-specific calibration offset and scale parameters were calculated.

To evaluate the performance of the logger, the following three metrics are analyzed for the voltage and current channels: measurement accuracy, noise floor and bandwidth. In addition to these metrics, the range switching time and burden voltage is analyzed for the current measurement circuit, as well as the input leakage current for the voltage channels. Again, the SMU was used to generate the stepwise voltage and current sweeps for the accuracy measurements. The characterization of the DC accuracy was performed 24 h after calibration of the device. The noise level was characterized under zero input condition, i.e., under short circuit conditions for the voltage channels and open circuit conditions for the current channels. For characterization of the channel bandwidth, a Keysight 33600A waveform generator generated a frequency sweep of a sine wave from 0 kHz to 50 kHz with amplitudes of 40 mA and 1 V. Finally, the range switching time and transient burden voltage were measured for current

steps from 0 mA to 500 mA using a Tektronix MSO4104B oscilloscope.

2.4.2 Current Measurement Performance

We start with the presentation of the current channel performance characteristics. The characterization was performed at a sample rate of 1 kSPS, if not specified otherwise.

DC Accuracy

The DC accuracy of the current channel was measured using a current sweep across the full measurement range for both, the low current measurement path based on the feedback ammeter circuit and the high current measurement path using a shunt ammeter circuit. Comparing the measured currents to the evaluated setpoints reveals a DC accuracy of $0.03\% + 4\text{ nA}$ for the low and $0.09\% + 3\text{ }\mu\text{A}$ for the high current range.

Noise Floor

The noise floor analysis for the open circuit current experiment at 1 kSPS for low and high range show a root mean square (RMS) noise level of 1.33 nA and 1.34 μA , respectively.

Bandwidth

The analog bandwidth was characterized with a frequency sweep using the signal generator. The -3 dB frequency at 64 kSPS is 9.5 kHz. For lower sample rates, the bandwidth corresponds to the one of the analog-to-digital converter (ADC) low-pass filter, resulting in a bandwidth of 262 Hz at 1 kSPS.




Figure 2.4: Analysis of the transient burden voltage for different current steps from 0 mA to 500 mA, performed at time $t = 0 \mu\text{s}$. The results demonstrates a fast switching within $< 1.4 \mu\text{s}$, and a maximum burden voltage during that transition of 430 mV.

Burden Voltage

The measurement of the current channel's burden voltage shows a linear behavior across the full measurement range with a maximum of 53 mV at 500 mA, corresponding to an input impedance of 106 m Ω .

Range Switching Time

The experiment to characterize the current step response shows the following results: within only 1.4 μs from detection, the MOSFETs complete switching from the low to the high current range. This shows that switching in the analog front-end is one order of magnitude faster than the delay between two consecutive ADC samples at the highest sample rate. During this short period, the burden voltage experiences a short peak. However, the characterization of the burden voltage shown in Figure 2.4 for different current steps confirms

that the burden voltage still stays within 185 mV for current steps up to 100 mA and within 430 mV even for the maximum current step to 500 mA.

2.4.3 Voltage Measurement Performance

The voltage channel characteristics are presented below. Again, the values were recorded at a sample rate of 1 kSPS, if not specified otherwise.

DC Accuracy

The accuracy of the voltage channels was evaluated using a voltage sweep across the full measurement range. The absolute voltage measurements are then compared to setpoints to analyze the DC accuracy. The results show an accuracy of $0.02\% + 13\ \mu\text{V}$.

Noise Floor

For characterization of the voltage channel's measurement noise, a zero input voltage experiment was used. The noise floor measurement results show an RMS noise of 5.9 μV .

Bandwidth

The analog bandwidth was analyzed using a frequency sweep from 0 Hz up to 50 kHz. The results show that the bandwidth matches the specification of the ADC low-pass filter: at 64 kSPS a $-3\ \text{dB}$ frequency of 10 kHz was measured. At a sample rate of 1 kSPS the bandwidth was reduced to 262 Hz.

Input Leakage

Minimal input leakage of the voltage channel is very important when performing current measurement in parallel to not

Table 2.1: Performance characteristics of the RocketLogger. Noise and bandwidth values are shown for 1 kSPS and (64 kSPS).

	Metric	Value/Range
Current ± 500 mA ($\times 2$)	Total dynamic range	172 dB
	Burden voltage at 500 mA	53 mV
	Noise high range	1.34 μ A RMS (50.6 μ A RMS)
	Low current range	± 2 mA
	Noise low range	1.33 nA RMS (276 nA RMS)
	Range switching time	1.4 μ s
	Transient burden voltage	max. 430 mV for < 1.4 μ s
	Measurement bandwidth	262 Hz (9.5 kHz)
	Accuracy low range (24 h)	0.03 % + 4 nA
Accuracy high range (24 h)	0.09 % + 3 μ A	
Voltage ± 5.5 V ($\times 4$)	Noise	5.9 μ V RMS (1.38 mV RMS)
	Input leakage	~ 5 pA
	Measurement bandwidth	262 Hz (10 kHz)
	Accuracy (24 h)	0.02 % + 13 μ V
Digital Input ($\times 6$)	Input leakage	< 1 pA
	Threshold voltage	-6 V to 6 V, configurable
General	Sample rate	1 kSPS up to 64 kSPS
	Data logged	up to 8.29 GB/h
	Logger dimensions	103 mm \times 68 mm

impact the current channel readings. Measurements with the power supply at the maximum input voltage of 5.5 V show a leakage current of around 5 pA, corresponding to an input impedance in the range of 1 T Ω .

2.4.4 Performance Summary and Positioning

Table 2.1 summarizes the performance characteristics of the RocketLogger at its base sample rate of 1 kSPS, with additional characteristics for the highest sample rate of 64 kSPS. With the shown accuracy, the very large dynamic measurement range and the portable design of only 103 mm \times 68 mm, the RocketLogger presents a well-balanced trade-off between measurement performance, portability and features. Com-

plementing these power measurements with digital state tracing and environmental monitoring, the logger satisfies all requirements for long-term in-situ measurement of energy harvesting system designs.

In terms of measurement performance and portability the RocketLogger is positioned in between high accuracy, stationary digital power meters and embedded power loggers, as the comparison of the logger with a representative solution of each category shows in Table 2.2. With some unique features like the integrated environmental logging that is not provided by other solutions, the RocketLogger presents a perfect tool to support the design of novel energy harvesting-driven devices.

2.5 Case Studies

Harvesting-based system design has many different aspects which require extensive measurement data. In following case-studies we highlight some of the main difficulties in obtaining high quality measurements and demonstrate how the unique features of the RocketLogger enhances this aspect in the design process of harvesting-driven systems. To this end, three different case studies are considered. The low-power duty-cycling optimization in Section 2.5.1 emphasizes the need for high accuracy measurements from an independent observer. A real-world evaluation of multi-source harvesting circuit designed for wearables in Section 2.5.2 shows the importance of joint logging of the power and environment for detailed system evaluation in the field. The last long-term system evaluation in Section 2.5.3 highlights the difficulty of experimentally validating energy harvesting devices without long-term in-situ measurements.

Table 2.2: The RocketLogger is positioned between stationary digital power meters and embedded power loggers.

	Keithley Digital Multimeter [Kei13a]	RocketLogger (this work)	Embedded Power Logger (Nemo [ZX13])
<i>Voltage range</i>	1 nV to 1 100 V	6 μ V to 5.5 V	N/A
<i>Voltage accuracy</i>	0.004 % + 1.2 μ V	0.02 % + 13 μ V	N/A
<i>Input resistance</i>	> 100 G Ω	\sim 1 T Ω	N/A
<i>Current range</i>	10 pA to 2.1 A	1 nA to 500 mA	0.1 μ A to 202 mA
<i>Current accuracy</i>	0.05 % + 1.2 nA	0.09 % + 4 nA	1.34 % average, 8 % max
<i>Sample rate</i>	up to 2 MSPS	up to 64 kSPS	up to 100 kSPS
<i>Burden voltage</i>	\leq 350 mV	\leq 53 mV	\leq 130 mV
<i>Form factor</i>	Stationary lab device	Portable data logger	Embedded in circuit design
<i>Power supply</i>	Wall plug required	Battery powered	Application circuit
<i>Remote control</i>	Measurement control	Full device access	N/A
<i>Environment monitoring</i>	Temperature only (via voltage)	Digital sensor bus for temperature, humidity, illuminance, pressure, etc.	N/A
<i>Extendability</i>	Closed, proprietary system	Any Linux compatible program	Limited (firmware modification required)
<i>Cost per channel</i>	\$ 4900 to \$ 7300	\sim \$ 50	\sim \$ 10




Figure 2.5: Comparison of the power traces for four alternative timed sleep implementations, sampled at 1 kSPS. The close-up on the right side shows the different sleep power levels and the varying wake-up delay of the three lowest power options.

The systems designs and evaluations in the following Chapters 3 to 6 present additional, more extensive use-cases for the RocketLogger. Furthermore, the multi-year tracing effort to characterize indoor solar harvesting presented in [SGT19] was enabled by the RocketLogger.

2.5.1 Low Power Optimization

In a common prototyping scenario, optimizing system sleep modes can be a trade-off between sleep power and accurate timing. This illustrates the need for ultra-low power measurement capabilities and fast, seamless range switching. In this example, we consider a 1 s sleep phase between task executions. In this experiment an MSP430 platform was used and active tasks were simulated by turning on two LEDs. Figure 2.5 shows the recorded power consumption as a function of time for four different implementations of this application scenario. The right plot shows a close-up

of the second task activation. The *busy idle* is the baseline configuration where system performs busy idling at the core frequency of 8 MHz between tasks. Three alternative implementation make use of the lowest supported sleep state: *high freq.* describes sleeping with a timer clocked at 1 MHz using the high frequency system clock, *low freq.* denotes the same timer but clocked at 9.6 kHz from a low frequency oscillator, and *watchdog* denotes the use of the watchdog timer to time the wake-up for the next task execution. These results show the trade-off between accurate sleep delay and ultra low power consumption of different clock sources. During system prototyping, the high dynamic range of the RocketLogger is essential for evaluating and verifying different low-power design decisions and trade-offs.

2.5.2 Multi-Source Wearable Harvesting

An important phase of the harvesting circuit design is to perform real-world measurements to characterize environmental conditions, harvested power and conversion efficiencies. For wearable devices, this is particularly difficult due to the highly time-variant nature of the scenario. Furthermore, in-situ characterization demands for a portable device to be carried by the user. In this use case, we examine a multi-harvester wristband consisting of solar cells and thermoelectric generators (TEGs) [MBS⁺16]. In the test scenario shown here, the user walks outside during a warm, sunny day and then enters a colder, dark indoor space. Figure 2.6 shows harvester measurements and environmental conditions, in the upper and lower plots, respectively. For the harvester measurements, it should be noted that the harvested TEG power is in the order of 100's μ W, while solar power is in the mW range. The right hand axis shows the operating voltages of the solar cells and the TEG. This result shows that the




Figure 2.6: The harvesting power levels and input voltages of a multi-source harvesting circuit combining solar and thermoelectric energy harvesting. The lower plot shows the ambient temperature and illuminance traces describing the time-varying harvesting scenario.

solar harvesting power dominates outdoors, while the TEG generates more power indoor, although at a lower power level. This data is very valuable for subsequent iterations of system modeling and analysis to optimize important system parameters like harvesting efficiency. Without independent, synchronized measurements of multiple voltage and current channels, as well as environmental conditions, this harvesting wristband could not be accurately measured or evaluated in the field.

2.5.3 Long-Term In-Situ System Evaluation

One of the final prototyping stages involves testing and validating the system in the field. This requires long-term measurements to validate correct and energy efficient operation over an extended period of time. Figure 2.7 shows a 7 hour slice of a day-long experiment to measure and validate a solar powered sensor node [GSS⁺17]. Its harvesting circuit slowly accumulates energy in a buffer capacitor, following the energy management principles introduced in Chapter 5. Depending on the voltage level of this capacitor, a microcontroller is triggered to perform different tasks. The node had already been tested to be functional, and was mounted on an office wall exposed to indirect natural and artificial light. Figure 2.7 shows the node's illuminance during the experiment in the lower plot, which closely correlates to the harvested solar power in the upper plot. It can be seen that some fault occurred during the last hour of the experiment, keeping the buffer voltage high and the input power oscillating. Close-ups of when this fault happened are shown in the plots on the right. Around $t = 5.97$ h, after more than 3700 successful task executions, the buffer crosses a threshold but the trace indicates that the node entered an undefined steady state. Closer investigation revealed that the triggering mechanism




Figure 2.7: Power and ambient trace of a long-term in-situ validation of an energy harvesting-driven sensing system. Only this long-term in-situ test revealed the timing fault at $t = 5.97$ h, since it is tied to the time-varying operating points of both, the source and load.

had, under rare and very specific circumstances, a timing fault. Device tests under standard laboratory conditions would not have revealed this fault, since it is tied to the time-varying operating points of both the source and load. Only with in-situ measurement over an extended period of time is a thorough validation of devices designed for long-term deployment possible.

2.6 Summary

In this chapter we addressed the need for portable and accurate measurements in the design of energy harvesting system. We presented the RocketLogger, a measurement device capable of accurately measuring power with a high dynamic range and logging environment conditions in long-term in-situ deployments. Combining a shunt resistor and a feedback ammeter circuit with seamless range switching, we tackled the challenge of measuring the widely varying power in harvesting and low-power systems, while minimizing the the impact of the measurements on the device under test. The four voltage and two current channels integrated in a handheld measurement device enable simultaneous observation of multiple energy flows like the harvested and consumed energy. Extensive characterization of the measurement performance demonstrated a high accuracy and minimal impact on the device under test. Three use-cases highlighted how the RocketLogger's unique feature set supports recording highly relevant and accurate experimental data.

3

Testbed for Environment and Application Emulation

Energy harvesting systems strongly depend on the non-deterministic environment and need to satisfy highly dynamic application requirements. Systematic and thorough evaluation of these systems demands for tools that can consistently reproduce these wide-ranging conditions. To this end, we introduce a testbed that reproduces varying thermal and radiation properties in a controlled environment and is capable of sinking current profiles in a time- and event-triggered manner. The coordinated control of these boundary conditions enables detailed evaluation, exploration and dimensioning of different energy management aspects in harvesting systems. By reproducing environmental properties at a higher rate, the testbed also allows to substantially shorten the time needed for experimental evaluations. This approach enables fast and consistent evaluation of energy harvesting systems with a wide coverage.

3.1 Introduction

Testbeds are an invaluable tool for evaluating systems under controlled conditions. In particular in contexts where the primary factors affecting the system performance exhibit significant non-determinism or variability, they enable controlled system evaluation and detailed analysis of the impact these factors have.

Energy Harvesting System Evaluation

In the domain of energy harvesting systems, jointly managing the energy extraction from a non-deterministic environment and satisfying the dynamic load requirements represents a major challenge. Consequently, energy management solutions addressing these challenges need to adapt to highly dynamic operating conditions to operate efficiently. This has also implications on their evaluation: for thorough characterization, analysis, dimensioning and validation of these systems, energy harvesting systems need to be evaluated with full coverage of the multi-dimensional and widely varying operating conditions. In addition to these requirement in terms of dynamic range and precision, the interactions of the testbed with the target system must be time-synchronized. Therefore, an appropriate coordination through a global control scheme is required. For a feedback based design flow to benefit from these evaluations, the conditions under which they are performed have to be consistent to guarantee comparability among them. Furthermore, rapid evaluation is required to provide timely feedback in the design flow. Consequently, the interactions between the testbed and the system under test need to be sped up, in terms of generation of the physical environment, consumption of energy, as well as in terms of measurements. In summary, a testbed for energy

harvesting systems has to address the issues related to the environment's reproducibility and wide-ranging electrical load properties, while providing rapid evaluation of multi-dimensional configurations.

Contributions

In this chapter we address these challenges in evaluating energy harvesting systems by introducing a testbed that provides controlled emulation of the environment's physical properties and the application requirements. Controlling these two major sources of non-determinism brings thorough evaluation into the design of reliable and efficient energy harvesting systems. The emulation of these two factors is complemented with centralized measurement and control to enable deliberate scaling of the time-domain for fast evaluation.

Specifically, we summarize our contributions as follows:

- We present two environment emulators that provide consistent reproduction of thermal and visible light properties of the environment. This enables in-depth analysis of thermoelectric and photovoltaic energy harvesting systems with full coverage of their harvesting conditions.
- We introduce a novel programmable current sink that allows time- and event-triggered replay of arbitrary electric load characteristics, thus supporting accurate reproduction of the dynamic application requirements of low-power systems.
- We integrate the environment and electric load emulation in a harvesting testbed. Adding centralized feedback control enables considerate speed-up of exhaustive evaluations of energy harvesting systems.

Leveraging the RocketLogger platform introduced in Chapter 2 that contributes comprehensive system observability, we integrate these different emulation, measurement, and control aspects into a powerful testbed system. Ultimately, the presented testbed solution provides fast and consistent evaluation with unprecedented coverage. In two case studies we highlight the significantly simplified and automated experimental evaluation in the design of energy harvesting systems.

Outline

The remainder of this chapter is structured as follows. Related work is covered in Section 3.2. We introduce the emulation of physical properties of the environment in Section 3.3. Subsequently, we detail the electrical load emulation in Section 3.4, before we discuss the integration of the emulators with flexible time control and measurements in a global testbed system in Section 3.5. The benefits of the testbed implementation are highlighted in two case studies in Section 3.6, before we summarize our contributions in Section 3.7.

3.2 Related Work

We cover previous work related to the emulation of the environment's physical properties and energy harvesting sources in general. Further, we review solutions presented in the area of reproducing electrical loads with dynamic behavior.

Environment Emulation

For photovoltaic and thermoelectric harvesting considered in this chapter, the focus lies on the emulation of temperature

gradients and solar radiation.

An approach to characterizing thermoelectric generators (TEGs) taken in [AGCCLS16] was to experimentally determine its electrical and thermal equivalent model parameters, and to investigate their relationship. Alternatively, the I-V characteristics of TEGs were extracted by varying their temperature gradient and connecting different power resistors as load [IAJCS17].

In the characterization of solar cell materials, accurate reproduction of the solar spectrum is key to perform low-level characterization of solar panels like measuring quantum efficiency [HD16]. To this end, solar simulators have been developed for large scale concentrator photo-voltaic systems [DAS08]. Similarly, highly accurate commercial sun spectrum sources exist for device testing [Sci20], but are very costly with a over 80 k USD for an illumination area of 45 cm × 45 cm. Lower cost alternatives using arrays of colored LEDs and a well characterized spectrum have been proposed by Bazzi et al. for testing solar cells [BKS⁺12].

Similar to previous work, we rely on dual sided temperature forcing to reproduce consistent temperature gradients for evaluating thermoelectric harvesting. Unlike the characterization testbeds for solar cell materials, we aim at emulating diverse solar harvesting scenarios with illuminance levels ranging from dimly lit indoor environments to direct sunlight.

Energy Harvesting Emulation

In the embedded system domain, spectrum dependencies are often abstracted away, as the electrical energy input to the system is the decisive factor. Instead, emphasis is put on either reproducing illuminance conditions [MEK⁺16, HRB⁺17] or electrical circuits exhibiting similar behavior [BBR⁺15]. Other works abstract away the physical environment and transduc-

ers and directly emulate their electrical behavior [HSS17, LFD⁺17, FdJ19], or control the input of harvesting circuits to force them to recorded operating points [GCZ19].

In this work we focus on emulating the physical properties of the environment as they present the original source of the non-determinism introduced into energy harvesting systems.

Electrical Load Emulation

Electronic current sinks are available in various forms. Bench-top source measure units (SMUs) [Kei19] can sink currents with high accuracy and are programmable.

There also several proposals of custom-built, configurable loads. Jian et al. designed an interleaved multi-phase pulse width modulation (PWM) scheme to mitigate ripples in power consumption [JWK⁺13]. In [Kaz07] the author presented a programmable DC load with different operation modes to emulate arbitrary load profiles. Upadhayay et al. developed a switching-converter-based electronic load with a high bandwidth [UMJ12].

While all the aforementioned works report high precision, they aimed at high power electronic loads. With focus on the embedded energy harvesting domain, we target emulation of currents ranging from a few nA to peak active currents in the 100's mA range.

3.3 Physical Environment Emulation

In this section, we first summarize the requirements for the physical environment emulation. We present two physical environment emulators afterwards that address these requirements for the evaluation of thermoelectric and photovoltaic harvesting, and we characterize their performance.




Figure 3.1: The architecture of the thermal testbed with two software controlled Peltier elements for consistent temperature gradient emulation. The adjustable distance enables evaluation of standalone thermoelectric generators (TEGs) or complete thermal systems including thermal guides.

3.3.1 Environment Emulation Requirements

The emulation of the environment’s physical properties needs to offer a high accuracy to provide consistency across experimental evaluations. This enables a comparison of the system performance under the same environmental conditions. The high variability of these properties demands the emulation of a wide range for full coverage of the dynamics in which the system operates. Lastly, these conditions should be adjustable in real-time to allow for an emulation with specific environmental traces and analysis of the system under realistic environmental dynamics.

3.3.2 Thermal Gradient Emulation

The thermal testbed architecture consists of two digitally controlled Peltier elements that offer cooling and heating of a metal plate terminal. Figure 3.1 illustrates the architecture of this testbed. We use commercial CP-061HT Peltier coolers




Figure 3.2: The thermal testbed implementation. In this example, a thermal harvesting system consisting of a thermoelectric generator (TEG) and an attached copper thermal guide is evaluated in between the two Peltier elements.

from TE Technology and the corresponding TC-720 unit for control. This unit integrates feedback control of the Peltier element's temperature and measures the temperature of the forced terminal. Setpoint configuration and measurement readout are exposed through an USB interface.

As shown in Figure 3.2, these two Peltier coolers are installed on a custom built frame with the temperature-controllable terminals facing each other. The adjustable distance between them allows inserting thermal harvesters of various sizes. This enables evaluation of standalone TEGs, as well as their integration into a thermal system including thermal guides.

The energy transfer is limited by that of the Peltier elements, which can extract/insert up to 63 W at a 0 °C temperature difference to the ambient. In combination with the external controller, this enables accurate and fast transitions




Figure 3.3: The solar testbed consists of a dimmable LED light source, a mirror box for a uniform illuminance distribution and a Peltier element to control the solar panel’s temperature. Temperature and illuminance sensors enable logging of the emulated conditions.

between temperature setpoint.

3.3.3 Visible Light Emulation

The architecture of the solar testbed introduced in the following is illustrated in Figure 3.3. We reflect on the most important design aspects of the custom-built solar testbed that enables consistent emulation of illuminance conditions ranging from dimly lit indoor environments to direct sunlight.

Light Source

The light source in the testbed consists of light-emitting diodes (LEDs). They are compact, cost effective and their illuminance is well controllable over a large range. In addition, their short warm-up time guarantees a fast and accurate response during emulation. The high efficiency of LEDs also simplifies heat management when emulating the high illuminance levels of direct sunlight. On the downside, LEDs

cannot accurately reproduce the spectrum of the sunlight. However, as was shown in [BKS⁺12], our approach could be extended to an array of colored LEDs to improve reproduction of the spectrum.

In the testbed Aventrix 2x2 matrices consisting of four NVSWE21A LEDs were employed as light source. Their high efficiency and luminous flux enable emulation of direct sun light conditions. A total of nine modules were mounted on a heat sink to cover a test area of 180 mm × 180 mm. To avoid high supply voltages, the control of these LED modules is distributed to four dimmer circuits that are discussed in more detail below.

Illuminance Uniformity

For consistent evaluation of solar panels it is important to illuminate them uniformly [SC08]. Our simulations of the illuminance distribution showed that a uniform illuminance distribution is achieved when the LED light sources are arranged in an evenly spaced grid and the light is reflected at the boundary of the light box. Consequently, we use mirrors to build a shielded light box. Illuminance measurements with only one of the mirrors removed confirmed that the illuminance differs up to 10 % across the light box, while the configuration with dual sided mirrors exhibited only minimal variability.

LED Illuminance Control

The illuminance of LEDs is controlled by either adjusting their supply current, or using a constant current source and pulse width modulation (PWM). There exist control modules for both approaches, but they are typically not made for high accuracy dimming and their precision is often not specified. Therefore, a custom LED driver was implemented



Figure 3.4: The fully integrated solar testbed. The top half represents the mirror-shielded light box with the LEDs mounted below the heat sink. The bottom of the light box consists of a Peltier controlled aluminum plate to maintain a constant temperature of the solar panel during the evaluations (hidden in the shadows). On the right side of the testbed the control circuitry is mounted: the LED dimmer with high dynamic range on top, the SmartLoad (presented in Section 3.4) on the bottom left, and the RocketLogger providing the measurement capabilities and hosting the testbed control (discussed in Section 3.5).




Figure 3.5: To provide control over a large illuminance range of 10 lx to 120 000 lx, a custom two-stage LED dimmer circuit was developed. It combines step-down LED regulation and pulse width modulation (PWM) dimming. Multiple of these channels are provided for controlling a total of 36 power LEDs.

to provide accurate control of a wide illuminance range. The dimmer circuit shown in Figure 3.5 combines a constant step-down current regulation using an LM3409HV buck controller and PWM dimming. To provide a high dynamic range, the step-down regulator is operated at two different currents, and PWM dimming is used for detailed illuminance control. Unfortunately, PWM dimming introduces ripples in the output current that could be captured by solar panels. To reduce this undesired effect to a minimum, we use a high modulation frequency and output filtering. The step down regulator's range and the PWM signal are controlled by a microcontroller that exposes an I²C interface for digital control of the illuminance level.

Temperature Control

As the solar panel performance is considerably affected by the ambient temperature [SP09], we include a CP-036HT Peltier element to control the temperature of the panel inside the

light box. The same TC-720 controller used in the thermal testbed is employed for precise temperature control and measurement. This guarantees consistent evaluation of solar harvesting when the LEDs emit significant amounts of heat at high illuminance levels, and enables temperature dependent performance analysis of the solar panel.

Environmental Condition Measurement

The emulated illuminance level and the temperature are monitored for the purpose of dimmer calibration and feedback on the emulation consistency when traces with fast transients are emulated (see also the discussion regarding emulation speed-up in Section 3.5.2). The temperature of the testbed plate is read out from the Peltier controller that measures this temperature for feedback control. Four TSL4531 ambient light sensors are placed at the edges of the illuminated test area to measure the illuminance. With a wide measurement range from 3 lx to 220 klx and a specified sensitivity of 1 lx, these sensors accurately cover the full illuminance range emulated in our testbed.

3.3.4 Solar Testbed Range and Accuracy

We evaluate the solar testbed's accuracy to validate that it satisfies the requirements of consistent emulation of a wide range of illuminance levels.

Prior to characterization, we calibrated the testbed in a two-step approach using the integrated light sensors as reference. First, the LED output illuminance E_V is translated into an LED current I_{LED} using a linear model. Subsequently, the duty cycle for each current channel/range combination is calibrated using a 3rd order model to match the current I_{LED} .

The high efficiency of the LEDs enables the emulation of

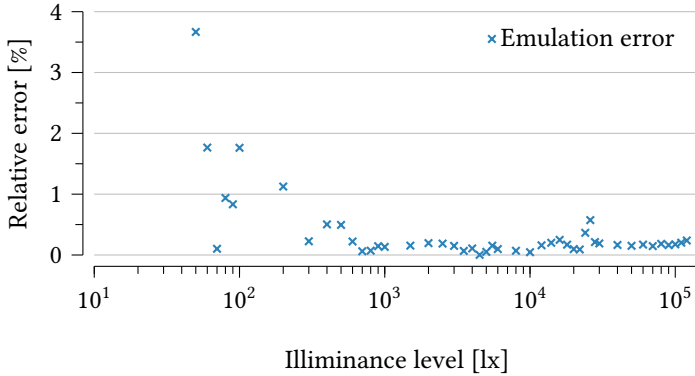


Figure 3.6: The illuminance emulation shows a high accuracy of $8.8 \text{ lx} + 0.5 \%$ over the range from 10 lx to 120 klx . A significant increase in emulation error below 100 lx shows the limitations of the presented testbed for extremely low illuminance levels.

illuminances from 10 lx up to 230 klx , which represents an illuminance higher than that direct sunlight (120 klx). With this coverage, photovoltaic harvesting scenarios with very high dynamic range can be emulated.

A few hours after calibration, the accuracy of the testbed is evaluated by sweeping the full illuminance range. Comparing the setpoints of the sweep to the sensor measurements results in the relative emulation errors shown in Figure 3.6. These results indeed confirm a high accuracy of the emulated illuminance of $8.8 \text{ lx} + 0.5 \%$. A noticeable increase in error is however reported for very low illuminance levels below 100 lx . This is predominantly attributed to the low PWM frequency that had to be chosen for the low illuminance range to prevent the step-down current regulator from overheating. In a redesign with stricter focus on thermal design, this can likely be resolved and the accuracy improved.

3.4 Electrical Load Emulation

This section introduces the SmartLoad, a programmable current sink that emulates the application's energy requirements. The stateful, time- and event-triggered emulation of current traces enables replaying of synthetic or recorded current consumption behavior of low-power systems.

3.4.1 Electrical Load Emulation Requirements

Similar to the environment emulation, the SmartLoad has to provide consistent emulation, while covering the full range of current consumptions in today's low-power systems that range from nA to 100's mA. The emulation requires output data rates and response times of typical embedded system in order to faithfully reproduce energy consumption traces. To employ the SmartLoad for transducer characterization, it has to support accurate current sinking at zero input voltage to allow for short circuit current characterizations.

3.4.2 SmartLoad Architecture

The architecture of the SmartLoad is illustrated in Figure 3.7. In the following, we cover the details of the analog current sink circuit and the digital control used for accurate replay of current traces.

Current Sink

The analog front-end uses a CSD18542KTT power MOSFET M_1 as controlled current sink. The feedback measurement from the shunt resistor R_{sh} is amplified by an AD8421 instrumentation amplifier. This signal is compared to the reference level set by an AD5761R digital-to-analog converter (DAC). The amplified signal difference provided by an OPA130




Figure 3.7: High-level schematic of the SmartLoad programmable current sink. A feedback controlled MOSFET sinks a current that is set by the digital-to-analog converter (DAC). The controller exposes an I²C interface for configuration and replays current traces in time- and event-triggered manner.

operational amplifier (op-amp) controls the MOSFET M_1 . For precise current control and to provide a high dynamic range, the shunt resistor value R_{sh} used for feedback measurement can be switched between three values. This enables more precise control using the DAC and at the same time keeps the burden voltage at the input terminals low.

An offset voltage of 1 V can be inserted into the current path to enable current sinking at zero input voltage. This is important for sinking from low voltage sources such as thermoelectric generators (TEGs) or to perform short circuit current characterization of transducers.

Control Loop Stability

To optimize the circuit for stable operation and highly accurate current control over the full operating range, we modeled the feedback control circuit. Using the device-specific transfer functions of the op-amp, MOSFET, and amplified shunt resistor measurement, the circuit's stability and output offset error were analyzed. A high feedback gain of $A_V = 100$ for the

linear amplifier was selected to reduce the offset error in the output current. To increase the very small stability margin at this amplification, a phase lag compensator was added in front of the power MOSFET to obtain a robust current control (not shown in Figure 3.7).

Event-Triggered Emulation

The digital-to-analog converter (DAC) is controlled by a microcontroller to provide stateful replay of stored current traces. Using the I²C interface, current traces are downloaded to the SmartLoad, and the replay of them is configured and controlled. In the present implementation, a replay is triggered through the control interface or external triggers. However, the functionality of the SmartLoad controller is extendable to support more complex state machines and state-dependent digital outputs. Such extensions allow for digital reconfiguration of the energy management dependent on the application state.

3.4.3 Electrical Load Evaluation

While the requirement of sinking short circuit currents at zero input voltage is guaranteed by design, we experimentally evaluate the accuracy and dynamic response behavior of the SmartLoad.

Prior to evaluating its performance, we calibrated the SmartLoad. A Keithley 2450 source measure unit (SMU) was used to source the current and to perform reference measurements. Current sweeps were performed for all three current ranges to determine their individual linear calibration parameters.

Table 3.1: DC performance of the SmartLoad.

Name	Range	Accuracy
High Range	5 mA to 500 mA	0.01 % + 80 μ A
Medium Range	50 μ A to 5 mA	0.008 % + 100 nA
Low Range	0 μ A to 50 μ A	0.003 % + 2.25 nA

DC Accuracy

Table 3.1 shows the DC performance of the SmartLoad 4 h after calibration. This evaluation was performed again using the SMU as source and for reference measurements at an input voltage of 0 V. The results show a very high accuracy of the sunk current for all three ranges. Higher input voltages likely result in a larger offset error in the lowest range, since the leakage current of the MOSFET will be in the range of several 10 nA.

Dynamic Behavior

Table 3.2 shows the worst case step response times of the SmartLoad when performing a full range current step. The values shown represent the delay from the time a new setpoint was set until the output has reached 90 % or 99 % of the step size. In summary, the maximum delay of the SmartLoad is less than 400 μ s if switching the current range. This performance is sufficient to faithfully emulate the behavior of typical low-power systems.

Input Voltage Step Rejection

The last parameter evaluated is the SmartLoad's behavior under input voltage steps from 0 V to 5 V forced by the SMU. The result of this experiment when sinking a current of 10 μ A at the lowest range is shown in Figure 3.8. Only very short

Table 3.2: Dynamic performance of the SmartLoad.

Range	% of Step Size	Time
High Range	90 %	32 μ s
	99 %	173 μ s
Medium Range	90 %	15 μ s
	99 %	138 μ s
Low Range	90 %	145 μ s
	99 %	340 μ s

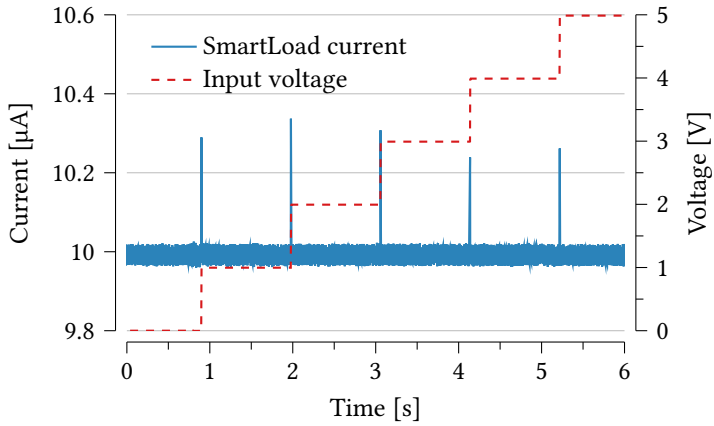


Figure 3.8: The input voltage step rejection evaluation of the SmartLoad shows very short current spikes of about $0.3 \mu\text{A}$ at the lowest current range.




Figure 3.9: Combining the different emulation tools with mixed-signal measurements and a central control enables a streamlined system emulation flow.

spikes of about $0.3 \mu\text{A}$ were observed in this current range, as the current sink's feedback loop has to adapt to the changed operating point. In the high current range these spikes were so small that they disappeared in the measurement noise of the SMU. This strong input voltage rejection is important to guarantee repeatable and accurate emulation of the current consumption under dynamic voltages.

3.5 Automated System Evaluation

To employ the environment and load emulators presented in the previous sections, they need to be integrated in an overall systems and controlled jointly. Global control enables automated and exhaustive evaluation of multi-dimensional configurations. Scaling of the time domain speeds up the evaluation, but demands careful consideration of potential side effects.

3.5.1 Automated Emulation Control

We combine the emulation of the physical environment and of electrical consumers with the mixed-signal tracing capabilities of the RocketLogger presented in Chapter 2. We leverage the extensible architecture of the logger and implement the mechanisms for central control of the emulation and measurements. Time-series inputs specify which environment parameters and current consumption to emulate, as is illustrated in Figure 3.9. During the experiments, various energy flows, environment conditions and system states are traced. This enables detailed analysis of the system performance and behavior of the energy harvesting system.

We focus on the two aspects of relaying stored traces for dynamic system evaluation and of exploring the system performance at predefined setpoints. However, the control infrastructure can be extended to emulate the conditions streamed from a remotely measured environment in real-time, for example.

3.5.2 Controlling the Time Domain

To accelerate the experimental evaluation, the time domain of the emulated traces can be scaled. Time scaling with a user-defined procedure is performed directly on the time-series inputs of the system. This flexibility allows to not only scale with a fixed speed-up factor, but also selectively using time-variant scaling methods.

The accelerated replay of these traces leads to much faster transients that need to be emulated. Therefore, it is crucial to consider that these transients can only be emulated within given bounds (see the step response characterizations in Sections 3.3 and 3.4). Developing a general methodology to determine the speed-up factor is difficult for many reasons:

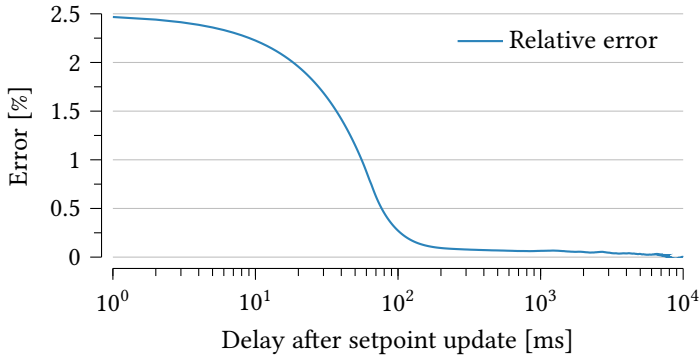


Figure 3.10: The setpoint characterization measurement of an AM-5412 solar cell shows drastically increased error if the analysis is performed using data points too close to the setpoint change.

the parameters of the system under test like thermal and electrical capacities, or fast transients in the recorded trace can result in faster dynamics than the testbed can emulate. This can significantly distort the experimental evaluation and may even lead to invalid conclusions. However, tracing the emulated conditions with the integrated sensors enables their comparison to input traces and allows to validate whether the dynamics of the trace were emulated successfully.

3.5.3 Rapid Setpoint Characterization

A typical task in system design is evaluating a large number of setpoints for design space exploration. Therefore, rapid setpoint characterization is important for high coverage and fast feedback. However, changing too rapidly between the setpoints can seriously distort the measurements due to transient effects. An analysis highlighting this issue for an AM-5412 solar cell characterization is shown in Figure 3.10: it shows the deviation of the measurements from the long-term

average within a 1000 ms window starting at the specified delay after the setpoint update. It reveals a significantly increased error if performing the analysis too early. To prevent this from happening, the testbed control provides a setpoint characterization mode. In this mode, the testbed control performs steady-state detection based on the algorithm presented in [KH13]. Using this feedback mechanism, the characterization is automatically sped-up in a controlled manner. A system designer only provides the list of setpoints instead of time-series data and the signal and analysis time window to use for steady-state detection. The time window is dependent on the system dynamics, but could be identified using a setpoint update experiment similar to the analysis shown in Figure 3.10.

We demonstrate this highly effective procedure in a transducer characterization case-study in Section 3.6.1.

3.6 Case Studies

In two case studies, we demonstrate how the testbed supports rapid evaluation in various aspects of the energy harvesting system design. First, we automatically characterize thermoelectric generators (TEGs) and solar cells. The second case study focuses on the evaluation of an energy-driven execution using an event-triggered application.

3.6.1 Automated Transducer Characterization

We demonstrate the automated characterization of the transducers used in harvesting systems. Thanks to the fully integrated harvesting testbed and its centralized control, this characterization requires very little effort: exposing the transducer to the desired operation conditions in the thermal

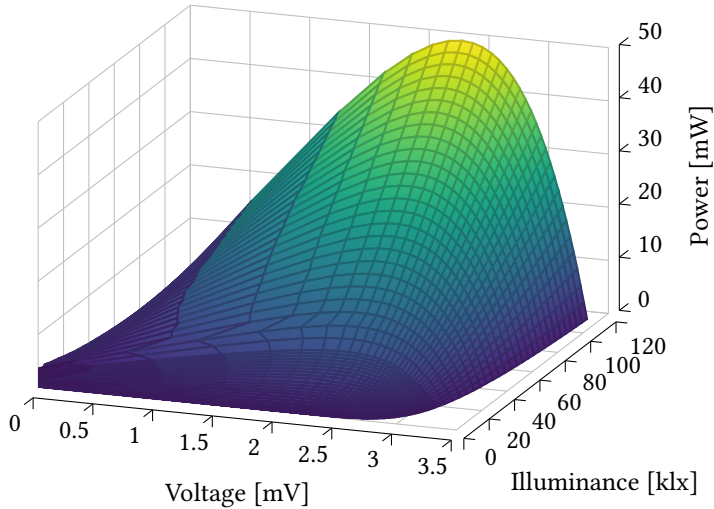


Figure 3.11: The automated transducer characterization enables rapid and exhaustive evaluation over a wide range of operating points. The illustrated characterization of an AM-5412 photovoltaic cell at a temperature of 25 °C consists of 2 160 setpoints that were recorded in only 76.6 min.

or solar testbed and directly connecting its output to the SmartLoad are all the necessary steps. After specifying the list of setpoints to evaluate and the required interval of steady input conditions, the testbed control handles sweep and steady state detection of these setpoints automatically.

Figure 3.11 shows an exhaustive multi-dimensional characterization of an AM-5412 solar cell, i.e., the transducers output power as a function of the illuminance and voltage levels of the emulated operating point. The panel was characterized for the full range from an illuminance level of 100 lx (dimly lit indoor) to 120 klx (direct sunlight), and from open-circuit to short-circuit conditions. The plot consists of a total of 2160




Figure 3.12: The thermal capacitances of the testbed and limited output power of the testbed demand for 30 h for a 461 setpoint characterization of a thermoelectric generator (TEG). The steady state detection guarantees accurate measurements without transient effects originating from setpoint switching.

characterization points that were recorded in less than 80 min.

Similarly, Figure 3.12 shows the characterization performed for a TEG241-150-29 TEG. Due to the thermal mass of the entire system and the limited output power of the Peltier element, the delay for reaching a stable setpoint measurement is much higher. For this reason the recording of the 461 setpoints took about 30 h. However, the fully automated process allows the characterization to run in parallel to other system design optimizations and guarantees accurate setpoint measurements.

These characterizations offer the exploration of many transducers and enable informed decision about the best performing option for a specific application scenario (see also

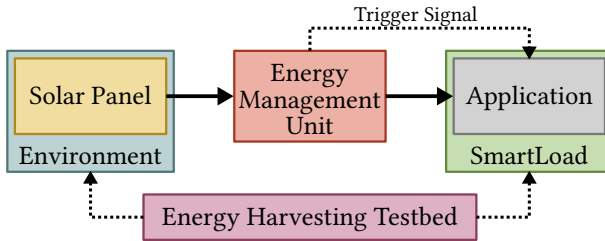


Figure 3.13: The experimental setup used for the energy management unit (EMU) evaluation. The SmartLoad replays the current trace previously recorded using the original application circuit.

Section 4.4). Furthermore, it supports determining system-specific parameters such as the calibration of maximum power point tracking (MPPT). Moreover, extensive validation at different operation points allows verification of existing transducer models or the development of new ones.

3.6.2 Energy-Triggered Application Execution

In the last case study, we demonstrate the use of the testbed in the end-to-end evaluation of an energy management system from the physical signal to the consumption of the energy by the application. For this purpose we use the energy management unit (EMU) to be introduced in Chapter 5, that provides the harvested energy in bursts in accordance with the application requirements. The setup of this case-study is summarized in Figure 3.13. An AM-5412 solar cell is used as transducer and connected to the input of the EMU. The solar cell is placed in the solar testbed, and the SmartLoad at the output emulates a simple demo current consumption trace of an LED flashing application. The trigger signal notifying the application of the availability of the next energy burst is connected to the SmartLoad for event-triggered replay of the

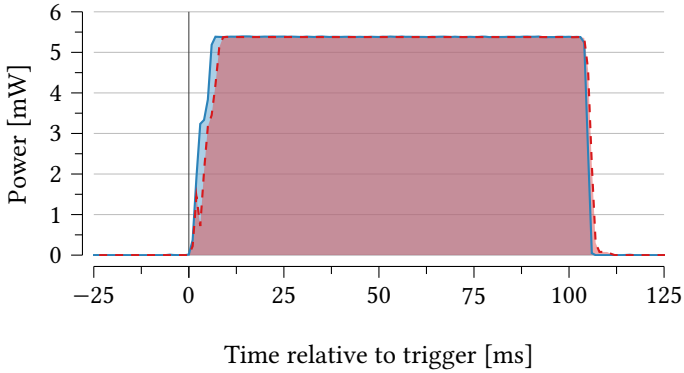


Figure 3.14: Comparison of the power traces of the original application and its emulation with the SmartLoad synchronized at the application trigger time.

application’s current consumption profile.

A detailed evaluation of the EMU performance under a wide range of harvesting conditions is given in Section 5.6. Here, we focus on the accuracy of the current trace replayed by the SmartLoad. Figure 3.14 shows the comparison of the replayed trace to the power consumption recorded with the actual application circuit. The emulated power closely follows the baseline application trace. Slightly slower dynamics in the emulated energy consumption are noticeable at the sharp transitions from and to sleep mode. With 0.8%, the relative error in terms of total consumed energy during the emulated application execution is minimal.

3.7 Summary

In this chapter we introduced a testbed for emulating the physical environment conditions and dynamic energy re-

quirements of the load. Emulation of the environment's thermal and visible light properties in combination with the capability to sink current traces in a time- and event-triggered manner provides the mechanisms for systematic analysis of energy harvesting systems. The consistent reproduction of the high dynamic range on the source and sink of energy enables evaluation with unprecedented coverage. Replaying traces at a higher rate offers substantially shorter evaluations and rapid feedback in the design of energy harvesting solutions. Two case studies highlighted the benefits of the testbed in exhaustive system evaluation.

The evaluation of the mechanisms presented in the subsequent chapters provide additional examples on how the testbed can be employed in the evaluation of energy harvesting systems.

4

Thermoelectric Energy Harvesting from Natural Temperature Gradients

We employ the design tools introduced in the previous chapters to develop a thermoelectric energy harvesting platform. Specifically, we exploit the naturally occurring temperature gradients at the interface between ground and ambient. To accomplish this objective, we address the challenges related to the low temperature gradients, thermal and electrical matching of thermoelectric generators (TEGs), and the efficient rectification of low voltages. The design and implementation of an efficient harvesting platform is possible as we approach the above mentioned challenges by an end-to-end modeling of the whole system, as well as extensive characterization and validation using the thermal testbed introduced in Chapter 3. Thorough validation and comparison with theoretical simulations of the harvesting system under

lab and real-world conditions attest a high accuracy of the developed model. Finally, we performed a multi-year deployment of the harvesting systems supplying a wireless sensor network (WSN) for environment monitoring. This extensive experiment supports our design choices and demonstrates a very high efficiency, with significant improvements compared to the state-of-the-art in both, the average and maximum harvested power.

4.1 Introduction

Energy harvesting presents an unprecedented opportunity to extend application lifetime [BASM16] and reduce maintenance cost, especially for long-term applications typically found in monitoring networks. In this chapter we focus on harvesting energy using thermoelectric generators (TEGs). TEGs convert heat flux directly into electric energy. They utilize thermal gradients originating from waste heat [DFG⁺09], the human body [TSM⁺17] or occurring in the natural environment [KWT⁺14]. More specifically, we investigate harvesting energy from gradients found in the ambient at the natural ground-to-air boundary, i.e., at the boundary of the atmosphere and the earth's surface. Short-term variations in air temperature are generally large compared to temperature changes in the outermost part of the earth surface or objects of the built environment [HGH11]. Radiation is the main driver of gradients found between ground, buildings or other solid objects and the atmosphere. During the day, especially under direct sunlight conditions, a surface exposed to radiation warms up. The radiative energy absorbed at the surface propagates into the material at a rate and with an attenuation depending on the material's thermal conductivity

resulting in a thermal gradient. The process reverses at night and is influenced by weather conditions, cloud cover, convection, etc. Therefore, thermal harvesting at the ground-to-air boundary exhibits seasonal patterns with a location dependent magnitude as well as daily patterns that can be exploited. Due to the bi-directional nature of the energy transport these patterns exhibit different temporal properties from other forms of energy harvesting, e.g. photovoltaics.

Alternative Approaches

Under direct sunlight conditions TEGs cannot compete with photovoltaic cells which have conversion efficiencies that are one order of magnitude higher [GDHE⁺20] than those of TEGs [YLYC18]. However, TEGs are symmetric, bipolar devices by design, and therefore capable of converting both directions of heat flux into electrical energy. Further advantages of TEG based energy harvesting are the utilization of a wide radiation spectrum, the tolerance for partial coverage (e.g. drop shadows, partial snow coverage) and the slow degradation of performance (aging, susceptibility to soiling) [Sny09]. This results in significant energy production at times when other approaches like photovoltaics fail, e.g., at night-time or at transition times between day and night. The nearly continual energy generation of a TEG harvesting based approach widens the design space to systems with only minimal energy storage. Such systems immediately use the harvested energy when it is produced, circumventing the losses due to further conversion and storage [BSBT14a]. Moreover, there are no moving or fragile (glass) parts, no acoustic emissions and the integration with a thermally conductive element and a radiator is simple, resulting in a rugged and highly reliable energy harvesting system.

Challenges

Contrary to thermal energy harvesting from process or body heat, where design and operation is governed by one heat flux direction, ambient thermal gradients are primarily driven by radiation that reverses its polarity at least once per day. This results in a bidirectional heat flux through the TEG module and consequently in bipolar output voltages. The small bidirectional temperature gradients of typically only a few Kelvin at the ground-to-air boundary present a veritably challenging harvesting scenario. Under these conditions a TEG generates small bipolar voltages in the range of several 10 mV to 100 mV. The voltage levels necessitate efficient low voltage rectification and voltage up-conversion to store the harvested energy in a buffer and/or supply it to an application circuit. To maximize the energy extracted, the thermal harvester system has to exploit the largest gradient possible at the ground-to-air boundary, while matching the TEG's thermal resistance for maximum heat flux and harvesting efficiency. Equivalently, the internal electrical resistance of a TEG has to be matched by the energy extraction circuitry for maximizing the electrical energy transfer.

Contributions

In the course of this chapter we make the following contributions to address the challenged discussed above:

- We present for the first time an end-to-end model of a thermoelectric harvesting architecture for the ground-to-air boundary. This enables use-case specific optimization for long-term autonomous operation and different power requirements.
- A novel low-power rectification circuit to rectify small bipolar voltages with minimal losses is designed. Ex-

perimental evaluation demonstrates its superior performance compared to other semiconductor and mechanical switching solutions.

- We perform in-depth validation of our model. The system components are evaluated individually in a controlled lab environment and the end-to-end model is validated by deploying the system in a concrete wall scenario.
- We provide extensive real-world performance evaluation using a multi-year deployed environment monitoring application. The system used for this study is dimensioned according to our model.

The optimized harvesting system consisting of a thermal guide and a 100 cm^2 radiator is deployed with a fully functional wireless sensor node in a long-term experimental study. The platform harvests up to 27.2 mW under direct solar exposure, while also providing up to 6.3 mW during night-time when emitting heat into the ambient. This enables self-sustainable operation of the environment monitoring application with a $550\text{ }\mu\text{W}$ average power footprint that not only senses environmental quantities but also participates in a multi-hop network and communicates the obtained information.

Outline

The remainder of the chapter is organized as follows: after covering the related work in Section 4.2, an overview of the harvesting scenario and method is given in Section 4.3. The building blocks of the thermal harvesting subsystem are modeled and validated in Section 4.4. Subsequently, Section 4.5 introduces and evaluates the low voltage rectification

circuit and completes the end-to-end model with the electrical subsystem. Section 4.6 presents a long-term evaluation of the harvesting system with an environment monitoring application. In Section 4.7 we summarize our contributions and findings.

4.2 Related Work

This section reviews the contributions made in the field of thermal energy harvesting from process heat or ambient temperature gradients and covers previous work addressing low voltage rectification.

Thermoelectric Harvesting from Process Heat

Numerous monitoring applications that harvest energy from process or waste heat using thermoelectric generators (TEGs) have been proposed. Dalola et al. monitored the temperature of hot pipes while harvesting energy from the large temperature gradient they produce [DFG⁺09]. In [WCW⁺13] the authors presented a wireless sensor node for building energy management applications that harvests from wall heaters. Thermal energy harvesting was combined with solar in [TP11] to increase the power harvested in indoor scenarios. In [MSG⁺19] TEGs were used to power the monitoring blade degradation of band saws. Radioisotope thermoelectric generators rely on radioactive decay of isotopes as heat source and are used as long-term energy supply in space applications [Row91], [Row06, Ch. 53-56]. Harvesting from human body heat has been integrated into wristbands [TSM⁺17], clothes [DQC⁺17] and headbands [Leo11] to supply smart watches and vital signs monitoring applications.

These applications consider aspects like thermal and

electrical matching that are relevant for the investigated scenario. In contrast to our scenario, they typically operate with unidirectional heat flux and at higher temperature gradients (with the exception of the human body scenario).

Thermoelectric Harvesting from Temperature Gradients in the Ambient

Previous work has also examined the exploitation of temperature gradients occurring in the ambient. In [Mik03] thermal heat flux in soil was analyzed and the authors concluded that thermoelectric harvesting is feasible at the ground-to-air boundary. Subsequently, they demonstrated in [AM03] that thermal harvesters embedded in the upper layers of soil generate an average power of 6 mW during the summer months using a TEG of undefined type and size. In [WD12] this scenario was modeled and thereafter experimentally evaluated. They reported an average power of $10.4 \mu\text{W}/\text{cm}^2$ before voltage rectification and conversion during the summer months. In [PKHW18], the gradients in the upper soil layers were monitored and through simulations it was concluded that a wireless sensor node can be supplied from only this energy source. Stevens et al. studied the theoretical optimal placement of a harvester at the ground-to-air boundary [Ste04]. In a later experimental study they harvested an average power of 1 mW with finned thermal guides of 3.8 cm diameter [Ste13]. By also using a finned heat sink [GSC⁺19] exploited the temperature gradient that occurs in railway tracks due to solar radiation. Meydbray et al. evaluated thermal energy harvesting from the surface-to-ambient gradient and reported an average harvested power of $5.7 \mu\text{W}/\text{cm}^2$ for a system with a ceramic plate on the ambient side [MSS05]. In [LS02] the authors performed an experimental study of the heat flux for a ground-to-air har-

vester with a thermal guide into the soil and a power transistor heat sink on the ambient side. Based on their observations they estimated a peak power of 0.4 mW. A similar study that combined a thermal guide reaching 20 cm into the ground with a TEG reports a peak power of $2.83 \mu\text{W}/\text{cm}^2$ [Lor03]. A convection dominated tunnel wall scenario was modeled and evaluated in [MEK⁺12]. Utilizing the optimized source and load matching, they reported 70 mJ of harvested electrical energy per day. Datta et al. focused on an asphalt surface to lower soil layer harvesting scenario, and reported up to $200 \mu\text{W}/\text{cm}^2$ harvested power around midday [DDP18]. Harvesting directly at the pavement surface, the impact of different TEG surface embedding options, surface colors and materials was experimentally evaluated in [KLYL17]. Instead of using the ground as a large thermal capacity, it is also possible to use phase changing materials. In [KWT⁺14, ABB⁺17] the fast changing ambient in an aircraft was used to generate an average power of 22 mW during an 80 min flight. Verma et al. harvested at the ambient to water storage boundary and used a fraction of the harvested energy to actively circulate hot water using a pump [VS19]. During the summer their systems generated an average output power of $341 \mu\text{W}/\text{cm}^2$ in an open environment.

Similar to some of the above mentioned work we use a thermal guide to contact lower ground layers. However on the ambient side we use a black body radiator to exploit the large solar radiation during the day and maximize the emission during night. Furthermore, we model the system end-to-end from ambient conditions incorporating meteorological data to the final wireless sensing application.

Efficient Rectification of Low Voltages

Although irrelevant for unidirectional heat flux, voltage rectification is crucial for the reversing heat flux in the ground-to-air scenario that generates bipolar output voltages. To rectify the voltage before up-conversion the DoupleTip platform [MCS12] uses a harvesting solution that integrates polarity switching into its conversion stage [Sal10]. However the platform has to be optimized at design time for a specific operating point. To enable the use of adaptive unipolar conversion circuits, active rectification is required as traditional diode based rectifications [HH15, Ch. 1.6] are not feasible for the low TEG voltages. The solution presented in [TWK⁺14] supports cold-starting with a depleted energy storage at the cost of higher quiescent current due to the system's demand for negative voltages. Similar self-powered rectification designs were optimized for efficiency and low quiescent current [PHMM11, HSS12], but require input voltages of several 100 mV.

An active bridge rectifier stage for rectification is part of all circuits. While others focused on cold-starting the rectification, we rely on a sufficiently large energy buffer to bridge energy unavailability, and focus on an ultra-low power control circuit that rectifies the low voltages of several mV with minimal power path losses, thus minimizing the overall power footprint.

4.3 Thermal Energy Harvesting System

The thermoelectric energy harvesting system introduced in this chapter extracts energy from the naturally occurring temperature gradients at the ground-to-air boundary. After specifying the harvesting scenario, an overview of the system

and its components is given.

4.3.1 Thermal Harvesting Scenario

The surface temperature of rocks and the built environment is governed by radiation and typically follows short-term variable signals [HGH11]. Because the adaption rate decreases with depth in materials with considerably high thermal capacitance, the strong thermal signal present at the surface is attenuated and delayed when propagating through the material. Consequently, a strong bidirectional temperature gradient is observable over the first few centimeters of a surface exposed to radiation that follows a (daily) recurring pattern. This is illustrated with the example of a surface temperature profile recorded in a rock wall environment in Figure 4.1. The mean annual temperature gradient between the ambient and a depth of 20 cm in the shown scenario is 0.08 K. Despite the system being in long-term thermal equilibrium, it is possible to harvest energy due to the above mentioned effects. The proposed harvester system adapts its harvesting polarity to the heat flux direction, thus increasing the mean temperature gradient exploitable for harvesting to 3.05 K.

4.3.2 Thermal Harvester Architecture

This work presents a thermal energy harvesting platform that exploits the above described temperature gradients using thermoelectric generators (TEGs). The harvesting platform consists of the following components, also illustrated in Figure 4.2: a TEG (c) is placed between an ambient facing black body radiator (a) that absorbs or emits thermal and solar radiation and a thermal guide (b) connecting the system to the ground at depth. The TEG transduces the resulting

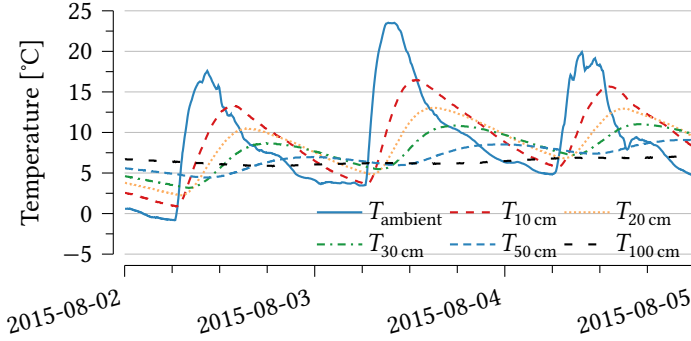


Figure 4.1: The rock temperature profiles from an alpine rock wall scenario illustrate the delayed and attenuated propagation of thermal energy through the rock material.

heat flux between air and ground into electrical energy. An electrical rectification circuit (d) reverses the polarity of the generator voltage according to the heat flux direction. The voltage converter and battery charge controller (e) converts the low TEG voltages to charge a battery and supply a wireless sensor node application (f).

4.4 Thermal System Model

The thermal harvesting platform is modeled end-to-end for design space exploration and system dimensioning. The model requires meteorological data as input and enables dimensioning a system that supports energy neutral operation of a wireless sensing application. Starting with the thermoelectric generator (TEG) transducing heat flux into electrical energy, the design and corresponding model of the thermal and radiation components are presented.




Figure 4.2: The architecture of the thermoelectric energy harvesting platform for the ground-to-air boundary: from the radiation regime on the ambient side (a), to the thermal guide connecting the ground to depth (b), using a TEG transducing the heat flux into electrical energy (c), and the electrical subsystem for rectification (d) and up-conversion of the low voltages to charge a battery (e) and supply the wireless sensor node application (f).

4.4.1 Thermoelectric Generator Characteristics

Exploiting the Seebeck effect, a TEG directly converts heat flux resulting from temperature gradients into electrical energy [Row06]. A TEG consists of p- and n-doped semiconductors with dissimilar thermoelectric properties that are electrically arranged in series and thermally in parallel. Applying a temperature gradient ΔT_{TEG} across the TEG leads to a heat flux through it, resulting in an open-circuit output voltage $V_{\text{TEG,OC}}$ proportional to the gradient ΔT_{TEG} .

$$V_{\text{TEG,OC}} = \alpha_{\text{TEG}} \cdot \Delta T_{\text{TEG}} \quad (4.1)$$

The Seebeck coefficient α_{TEG} of a TEG, in units of V/K, is a property of the type and shape of the semiconductor material. For temperature gradients up to several tens of Kelvins, the electrical properties of a TEG can be modeled as a temperature dependent voltage source as in (4.1) with constant internal resistance R_{TEG} [LMBAG10]. Consequently, the maximum power is extracted from the TEG when the resistance of the connected load matches the TEG's internal resistance, i.e., $R_{\text{load}} = R_{\text{TEG}}$. The resulting maximum harvested power depends quadratic on the temperature gradient across the TEG:

$$P_{\text{max}} = \frac{(V_{\text{TEG,OC}})^2}{4 \cdot R_{\text{TEG}}} = \frac{(\alpha_{\text{TEG}})^2}{4 \cdot R_{\text{TEG}}} \cdot (\Delta T_{\text{TEG}})^2 \quad (4.2)$$

Testbed Characterization of TEGs

The quadratic behavior of TEGs power for small temperature gradients was verified using the thermal testbed introduced in Chapter 3. TEGs were placed in between the two controlled Peltier devices and the testbed automatically evaluated the TEG output voltage and current at a wide range of thermal and electrical operating points. The exhaustive characterization of numerous commercially available TEGs for gradients up

to 10 K confirmed the quadratic output power model with negligible error for the evaluated operating points. This detailed exploration of different TEGs provides the foundation for an informed decision when selecting the best suited TEG for a specific system design.

4.4.2 Thermal Model of the Harvesting System

The thermal behavior of the harvester is approximated with a lumped parameter model using thermal resistances. It is based on the energy conservation law and assumes steady state conditions and an ideally insulated thermal guide. The thermal resistances of the components are determined by their respective materials and geometries. A two-dimensional (2-D) model, as illustrated in Figure 4.3, was selected, was selected to abstract the system and to keep the model complexity and number of parameters manageable. The inclusion of rock and thermal guide capacities, 3-dimensional geometries, and detailed contact properties with and without glue is expected to further improve the model accuracy, but is beyond the scope of this work. The temperatures at the system boundaries, T_{ground} and T_{radiator} , are either measured values or calculated based on the radiation model introduced in the following section. As a result the temperature gradient across the TEG module is calculated as:

$$\Delta T_{\text{TEG}} = \frac{(T_{\text{radiator}} - T_{\text{ground}}) \cdot K_{\text{TEG}}}{K_{\text{ambient}} + K_{\text{ground}} + K_{\text{TEG}}}, \quad (4.3)$$

where the thermal resistance on the ambient side is $K_{\text{ambient}} = K_{\text{radiator}} + K_{\text{contact}}$, on the ground side is $K_{\text{ground}} = K_{\text{rock}} + K_{\text{guide}} + K_{\text{transition}} + K_{\text{contact}}$, and K_{TEG} is the thermal resistance of the TEG. Equivalent to the electrical matching, thermal matching, i.e., $K_{\text{TEG}} = K_{\text{ambient}} + K_{\text{ground}}$, maximizes heat flux through the TEG for a given overall temperature gradient


[Sny09], resulting in maximum electrical output power (4.2). The presented thermal model is generic and allows using more advanced thermal designs, e.g. heat pipes or a finned design of the thermal guide or a heat sink replacing the black body radiator. Using such an advanced thermal design allows further tuning of the system for a specific scenario and offers the potential for higher harvesting efficiency. In this work the focus is on a simple and highly robust design for deployments in harsh outdoor environments.

Testbed Verification of the Thermal Model


The thermal harvester model was experimentally verified using the thermal testbed introduced in Chapter 3. For this purpose we used the thermal system dimensioned for the case study detailed in Section 4.6. Forcing of the system boundary temperatures T_{ground} and T_{radiator} allows evaluation of specific operating points and emulation of real-world traces as shown in Figure 4.1. As the contact resistance to the rock is not easily integrated and simulated in the testbed, the temperature at the end of the thermal guide is forced directly, therefore corresponding to a $K_{\text{rock}} = 0$ in the model. The comparison of the experimental to the simulation based TEG open-circuit voltage $V_{\text{TEG,OC}}$ in Figure 4.4 shows a close match of modeled and experimental behavior.

4.4.3 Radiation Model

The thermal system model is completed by integrating thermal radiation on the ambient side and consequently incorporating meteorological data into the model. The radiation model is based on energy conservation, dictating that the incident (short-wave) solar radiation S_{in} , reflected solar radiation S_{out} , absorbed (long-wave) thermal radiation L_{in} , emitted



(a) Mechanical structure of the thermal harvester system (not to scale).



(b) Lumped parameter model used to abstract the thermal guide.

Figure 4.3: The components of the lumped parameter model and their corresponding mechanical parts in the thermal guide design. The figure shows the dimensions resulting from the design choices made for the case study detailed in Section 4.6.

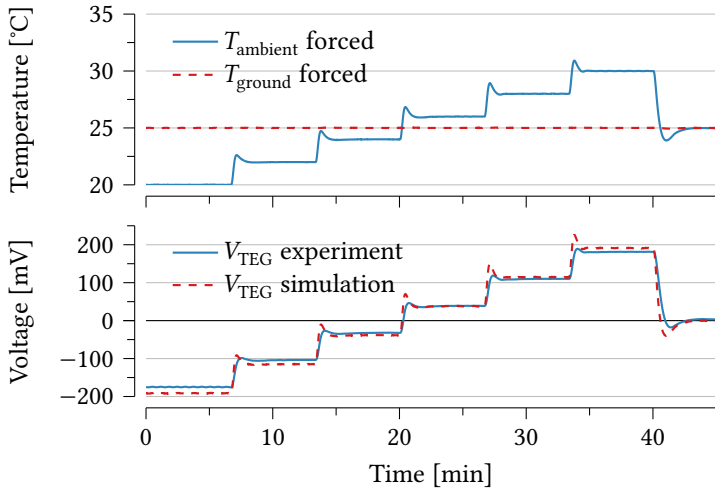


Figure 4.4: Comparison of testbed experiments and simulations for the thermal guide model, including the characterized TEG properties. The overshoot of the forced temperature is attributed to the step response of the control circuit in Peltier devices.




Figure 4.5: A representative two days excerpt of the model verification experiment in an urban concrete wall deployment. The measurements show the observed model inputs for the *thermal guide only* and the *end-to-end model* on the top. The comparisons of the models in terms of the TEG gradient ΔT_{TEG} in the middle and the matched output voltage V_{TEG} in the lower plot attest a high accuracy of the end-to-end model and its thermal guide component.

thermal radiation L_{out} , and heat flux through the thermal harvester into the ground Q_{ground} have to be balanced:

$$S_{\text{in}} + S_{\text{out}} + L_{\text{in}} + L_{\text{out}} + Q_{\text{ground}} \equiv 0 \quad (4.4)$$

The effective incident solar radiation S_{in} and reflected solar radiation S_{out} are derived from the solar radiation corrected by a multiplicative factor $\lambda \leq 1$. This correction factor accounts for the time-varying and deployment specific reduction of the radiation and includes the azimuth, elevation, altitude and visible horizon. For our deployments this factor is determined experimentally. However, for high precision models λ needs to be calculated dynamically every few minutes based on celestial data (sun), the true measured horizon, shading, as well as surrounding reflection parameters (albedo) [GPH⁺03]. The thermal radiation L_{in} depends on the air and ground temperature, T_{air} respectively T_{ground} , relative humidity RH , and cloud cover factor φ [Bru75, SB93]. The emitted thermal radiation L_{out} is a function of the radiator's temperature T_{radiator} and emissivity $\varepsilon_{\text{radiator}}$. Lastly, Q_{ground} incorporates the thermal model of the harvester which is a function of the harvester's overall thermal resistance K_{tot} and the temperatures T_{ground} and T_{radiator} (Figure 4.3). Because no closed form solution for T_{radiator} exists, L_{out} is first computed iteratively using a min search algorithm. Then, the radiator temperature is derived using the Stefan-Boltzmann law [GM74]:

$$T_{\text{radiator}} = \left[\frac{L_{\text{out}}}{A_{\text{radiator}} \cdot \varepsilon_{\text{radiator}} \cdot \sigma} \right]^{1/4} - 273.15 \quad [^{\circ}\text{C}], \quad (4.5)$$

where $\sigma = 5.67 \times 10^{-8} \text{ W m}^{-2} \text{ K}^{-4}$ is the Stefan-Boltzmann constant, A_{radiator} the radiator's surface area and L_{out} the emitted thermal radiation. The radiator temperature T_{radiator} is utilized to compute the resulting temperature gradient

ΔT_{TEG} according to (4.3). Subsequently, the TEG output voltage and its maximum electrical power are derived using (4.1) and (4.2).

Model Verification in Concrete Wall Scenario

To validate the full thermal model, the thermal harvesting system, dimensioned for the subsequently examined case study (see Section 4.6), was deployed in an urban concrete wall environment. Two independent harvesters were setup in the southwest (SW) facing artificial retaining wall. The deployed sensor nodes were configured to monitor the temperatures of the wall T_{ground} , the radiator side of the TEG $T_{\text{TEG,ambient}}$ and the thermal guide side of the TEG $T_{\text{TEG,ground}}$. The measured temperature T_{ground} in combination with the incident solar radiation and cloud coverage φ recorded by the close by (900 m) government meteo station enables the simulation of the radiator temperature T_{radiator} and TEG gradient ΔT_{TEG} . We use the same configuration as in the thermal guide testbed verification for the simulation, except that for contact of the thermal guide to the wall using concrete results in $K_{\text{rock}} = 1.25$. The simulation and experiment results are presented in Figure 4.5: the thermal radiator temperature is shown on the top and the TEG temperature gradient ΔT_{TEG} in the middle. The graphs demonstrate a good model accuracy for both the radiator temperature T_{radiator} and TEG gradient ΔT_{TEG} . Over the 43 day experiment period their respective mean absolute errors are 5.7 °C and 0.88 K. The errors can predominantly be attributed to the constrained horizon impacting radiation during morning and evening hours and the drop shadow at approximately 11:00. The bottom plot shows the experimentally observed and modeled optimal voltage at which the TEG is operating, which are discussed in the following sections.

4.5 Rectification and Electrical System Model

A novel low power circuit for rectifying the low voltages of thermoelectric generators (TEGs) is introduced and evaluated in detail. Subsequently, the remaining electrical building blocks of voltage conversion and application are discussed to complete the end-to-end model of the harvesting system.

4.5.1 Low-Power/Low-Voltage Rectification

Because the output voltage of the TEG depends on the heat flux direction (4.1), the bipolar harvesting scenario results in a bipolar output voltage. This demands for voltage rectification prior to supplying an application circuit. A low resistance rectification solution is key to minimize the losses in the power path. In addition, the low resistance is necessary to enable electrical matching to the TEG's low internal resistance and thus to maximize the harvested power.

Passive rectification circuits based on diodes are not a viable solution due to their forward voltage drop on the order of or higher than the TEG's output voltage. Consequently, an actively switched solution with minimal power footprint and power path loss is required. Solutions based on electromechanical switches, commercial solid state single pole, double throw (SPDT) switches and a custom designed MOSFET rectifier bridge are examined and compared with respect to their power path resistance and average power requirement. A control circuit consisting of two TLV3691 nano-power comparators is used for all three approaches, where one comparator observes the polarity of the input voltage and the other generates the inverted control signal.

Despite the considerable switching overhead of electromechanical switches, they provide a competitive solution




Figure 4.6: The low-voltage rectification circuit is based on a bridge of four enhancement mode nMOSFETs and a low-power control circuit.

since the polarity is typically switched only twice a day. The relays as well as the SPDT switches alter which terminal of the TEG is connected to which rectifier output depending on the voltage polarity. For the custom rectification circuit a full active bridge rectifier consisting of four n-type enhancement-mode metal-oxide-semiconductor field-effect-transistors (MOSFETs) is built as illustrated in Figure 4.6. The rectification does not require negative voltage control signals, because the TEG voltage remains at a low level compared to the control circuit supply. This allows keeping the quiescent current draw of the control circuit with 287 nA at a very low level. As a result, the circuit has an ultra low power footprint. Short-term SPICE simulations of the transient behavior of this solution, shown in Figure 4.7, confirm the correct operation of the proposed rectifier. Despite the very low switching frequency of typically two times a day, timely switching in the order of seconds is important to adapt to input-voltage inversion for continuous harvesting.




Figure 4.7: The short-term SPICE simulation confirms the functional correctness of the proposed MOSFET-based low-power low-voltage rectification.

Evaluating Alternative Rectifier Circuits

For real-world performance comparison, all three approaches are evaluated experimentally. The results of these experiments are summarized in Table 4.1. The large resistance measured for the SPDT switches leads to considerable power path losses. The proposed MOSFET and the relay solution demonstrate comparable power path resistances. In addition, the MOSFET design has negligible switching power. Although the relays show good performance for the considered metrics, they have significant drawbacks such as mechanical aging, shock sensitivity, unknown properties for very low voltage/current conduction and therefore unknown contact resistance characteristics in long-term operation. Therefore, the proposed MOSFET circuit is the overall best performing solution.

4.5.2 Electrical Load Matching and Voltage Conversion

The impedance of a harvesting circuit has to match the internal resistance R_{TEG} of a TEG to extract energy at the

Table 4.1: Performance Comparison of Different Switching Approaches for Low-Power Low-Voltage Rectification.

Method (Type)	$R_{\text{ON,TEG}^+}$	$R_{\text{ON,TEG}^-}$	avg. P_q
MOSFET (SiR404DP)	0.22 Ω	0.21 Ω	-
Relay (TXS2-L2-3V)	0.20 Ω	0.20 Ω	23 nW
SPDT (TS3A24159)	0.71 Ω	0.67 Ω	10 nW

maximum power point P_{max} (4.2). The TEG's internal resistance R_{TEG} can, with high accuracy, be assumed constant for low temperature gradients [TP11]. Therefore, the impedance can be matched during the system design phase, or adapted dynamically at run time using maximum power point tracking (MPPT). In the former case, the input impedance of the circuit is matched to R_{TEG} . For the latter case the input impedance is dynamically adjusted such that the input voltage $V_{\text{in}} = \frac{1}{2}V_{\text{TEG,OC}}$, guaranteeing operation at the maximum power transfer point.

The low temperature gradients of only a few Kelvins, as outlined in Section 4.3.1, lead to input voltages of a few 10 mV to several 100 mV. This voltage needs to be up-converted in order to charge an energy storage or supply a wireless sensor node application.

Commercially available harvesting management circuits combine these two important aspects and often also incorporate battery charge logic and output voltage regulation. Current solutions for TEG harvesting tend to either be passively controlled coupled inductor converters or actively controlled single inductor circuits [TSM⁺17]. The former enables harvesting from voltages as low as 20 mV. But their efficiency is limited due to the fixed voltage conversion ratio and internal linear voltage down regulation. The latter solution features a dynamic conversion ratio and input impedance matching

using MPPT but demands a comparably high minimal input voltage. The bq255xx harvesting management series [Tex13] are based on this principle and implement an ultra-low power control circuit to provide high efficiency at low current draw. Experimental evaluation shows that this particular architecture requires a minimal input voltage for harvesting of at least 60 mV.

For the considered harvesting scenario the bq25570 was selected due to its consistently high efficiency for a wide range of input voltages. The losses incurred by the higher start-up voltage are accepted as a trade-off for the high efficiency for larger input voltages. The bq25570 does not support negative input voltages and hence the TEG output voltages need to be rectified before serving as inputs for the harvesting circuit. In the bottom plot of Figure 4.5 the experimentally observed voltage at which power is extracted from the TEG is compared to the voltage at the theoretical maximum power point. The theoretical values of the optimal V_{TEG} are calculated once from the observed temperature gradient ΔT_{TEG} using solely the TEG model (4.1), and once from the meteorological data using the end-to-end model. The close match of these curves attests optimal load impedance matching and confirms the previously observed high model accuracy from radiation to the electrical energy. The irregularly observed spikes relate to the power point tracking mechanism of the harvester circuit.

4.5.3 Wireless Sensor Node Application

An application that performs local sensing and maintains a wireless network to forward recorded data is supplied with the harvested power. It is abstracted as an electrical load that performs duty-cycling, a concept widely adopted in low-power system design. The energy consumption of the sensor node is modeled by its power consumed during the active

and sleep states and the duty-cycle at which the application operates.

4.6 Thermal Harvesting System for Environment Monitoring Application

For demonstration of the applicability and evaluation of the proposed harvesting architecture the use-case of a wireless sensor node for long-term autonomous environment monitoring in steep rock wall is considered [BGH⁺09]. Specifically, the node monitors itself and senses environmental quantities like temperature profiles at varying depth, movement, thermal and electrical conductivity of rocks, and water pressure inside rocks. In parallel to local sensing the node participates in a multi-hop network and communicates the sensed information and system information like battery level and power consumption through this network. The average power footprint for operating such a sensor node is measured to be 550 μW .

The end-to-end model detailed Sections 4.4 and 4.5 is employed to dimension and implement a harvesting architecture optimized for the considered use-case. Analysis of the specific scenario reveals that the exploitable temperature gradient increases with depth from the ground surface, as well as with the surface area of the radiator. However, the system size is strongly restricted by deployment specific boundary conditions such as mechanical mounting and drilling equipment. In the following the details of the dimensioned system are discussed, before extensively assessing its performance in a long-term real-world deployment.

4.6.1 System Dimensioning and Integration

The resulting components and their dimensions for the thermoelectric harvesting systems are presented in the following.

Thermoelectric Generator (TEG)

A broad range of commercial thermoelectric generators (TEGs) were experimentally characterized. Characterizing their performance for low temperature gradients up to 10 K using the thermal testbed introduced in Chapter 3 allowed detailed exploration of alternative TEG modules. Finally, this led to the selection of a Thermalforce 241-150-29 TEG (of 30 mm × 30 mm) with a thermal resistance of 2.27 K/W and a large Seebeck coefficient of $\alpha_{\text{TEG}} = 111 \text{ mV/K}$, which together with the thermal system built around it provides the best harvesting performance.

Thermal Harvester

The length of the thermal guide is designed to be 200 mm, the maximum length permitted by the boundary conditions. This allows contact to the rock at depth to exploit high temperature gradients. Two diameters with values up to 25 mm are utilized to increase the thermal conductivity of the thermal guide. The thermal guide is made of copper due to its excellent thermal conductivity. A copper plate connects the thermal guide with the TEG. The other side of the TEG is in contact with the radiator. A black powder coated aluminum block is used as an approximation of an ideal black body radiator. Furthermore, the radiator is designed to have a large surface area of 100 mm × 100 mm to absorb the enough radiation to sustain the operation of the 550 μW load with sufficient margin. The radiator was designed with a thickness of 30 mm to also provide the housing for the TEG. The resulting

thermal system is of low complexity, mechanically robust, as well as shock proof, thus fulfilling the requirements imposed by the application's harsh environment.

Electrical Subsystem

The custom rectification and bq25570 harvesting circuit are integrated into the wireless sensor node. The extracted power is stored in a Saft MP174565 lithium-ion battery. These cells can be charged down to $-30\text{ }^{\circ}\text{C}$ as is necessary for our use-case and allow ample storage to bridge extended periods of energy scarcity.

Application Model and Integration

The considered application is modeled as an electrical load with an average power footprint of $550\text{ }\mu\text{W}$ [BGH⁺09]. The application is integrated to be solely powered from the harvested thermal energy. The sensor node is configured to monitor the performance of the harvesting platform. In the following we present the results of these in-situ measurements.

An exploded view of the system designed for the rock wall use-case, as well as a photograph of its deployment is shown in Figure 4.8.

4.6.2 Real-World Harvesting Performance

The final system was installed for long-term testing in a field site for high-alpine environmental research at 3 500 m a.s.l. [HGH11]. Two systems were deployed on a steep southeast (SE) facing and one on a near-vertical northwest (NW) facing rock wall. The system on the NW facing wall, position 31, and one on the SE facing wall, position 29, both had a thermal guide diameter of 16 mm, equal to the one in the




Figure 4.8: Left: exploded view of the harvesting system components consisting of the thermal harvester, TEG, rectification circuitry, and the wireless sensor node. Right: the system deployed in the rock wall for the long-term case study (photograph of the SE position 29).

urban concrete wall experiment. The second system in the SE facing location, position 28, had a larger diameter of 25 mm. Figure 4.8 illustrates the deployment of the SE facing position 29.

Overview

For each position between 6 and 21 months of data were captured. The average harvested power was 1.21 mW, 1.09 mW, and 0.76 mW for positions 28, 29, and 31, respectively. Figure 4.9 shows the evolution of the battery voltage relative to the marked starting point in the upper part. The lower part depicts the difference of the battery voltage over 3 days to amplify and visualize the momentary dynamics throughout the year. The curves are wrapped around at the end of the year. Periods with steadily decreasing ΔV_{bat} are longer bad-weather periods typically coinciding with a local increase in snow cover. Differences in snow cover result in position 28 starting to generate energy at the beginning of the year before position 29 does despite their proximity. The battery voltage for all positions increases during the summer months because of the longer periods of direct incident solar radiation. This effect is more pronounced on the SE than the NW facing deployments. Nevertheless, the NW facing node sustains energy neutral operation during the summer. The harvested energy however is not sufficient to compensate for the harvesting deficit during the winter.

Detail Excerpt

The 10 day excerpt from spring 2017 in Figure 4.10 shows the temperature gradient ΔT_{TEG} across the generator and the total harvested power P_{harv} . Initially, snow cover insulates all three positions resulting in an almost zero ΔT_{TEG} . As the snow cover decreases position 28 starts harvesting more energy on

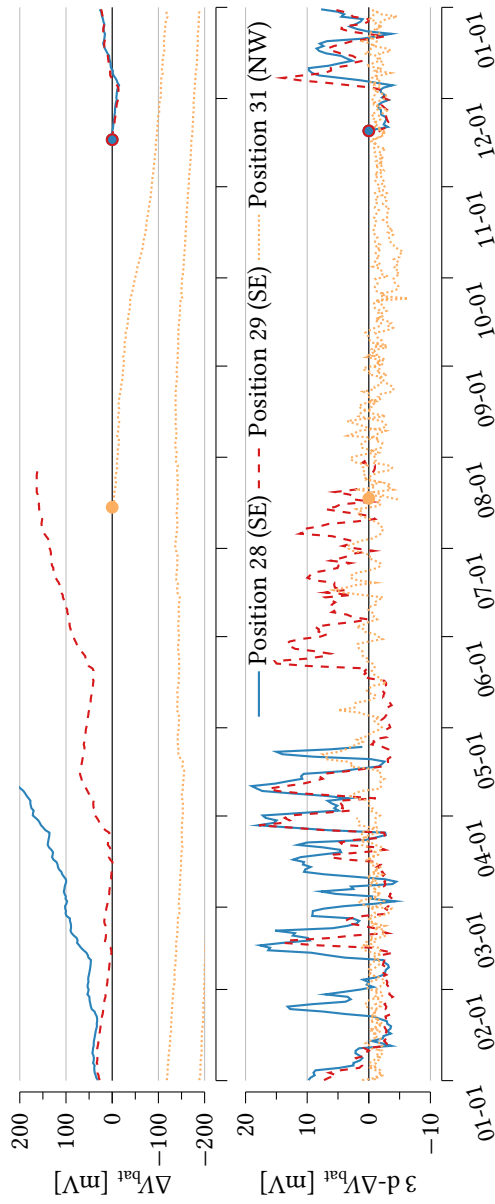


Figure 4.9: Long- and short-term (3 day differences) battery voltage evolution showing steady charging at the south locations.




Figure 4.10: Ten days of temperature difference and the resulting harvesting power showing the influence of weather and location.

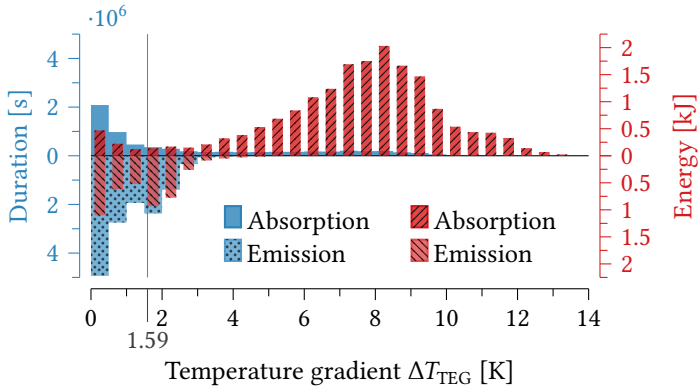


Figure 4.11: Histogram of the temperature gradients and harvested energy during the 237 days case-study of the SE facing sensor node deployment at position 29.

March 19. Position 29 follows the next day. The variability during March 21 indicates variable weather with intermittent cloud cover. During the following four days the nodes are not exposed to direct incident solar radiation because of constant cloud cover. Nonetheless, ΔT_{TEG} continues to follow a daily pattern. On March 26, a perfectly sunny day, the SE facing positions 28 and 29 generated a peak power of 25.4 mW and 26.5 mW, respectively. Because the solar elevation angle increases in springtime, position 31 starts receiving direct sunlight in the late afternoon which the temperature trace reflects. It must be noted that during the time window of the detail excerpt, the harvesting circuit of position 31 malfunctioned resulting in significantly reduced harvested power. However, the temperature measurements show correct values, allowing to draw the above conclusions.




Figure 4.12: Histogram of the temperature gradients and harvested energy during the 625 days case-study of the NW facing sensor node deployment.

Operation Mode Analysis

Lastly, the distribution of the temperature gradients and the generated energy on the SE and NW facing deployments is analyzed. The histograms show the duration the harvester was operating at a given TEG temperature gradient ΔT_{TEG} and the total energy generated during these periods, Figure 4.11 for position 29 and Figure 4.12 for position 31. As expected, both systems are in a thermal equilibrium with a mean ΔT_{TEG} close to zero. However, the harvesting relevant mean absolute gradients are 1.59 K and 0.96 K for positions 29 and 31. The difference between both positions is mainly attributed to the frequency with which high ΔT_{TEG} , resulting from direct incident solar radiation, occur. The SE facing node harvests most of its energy during periods of high temperature gradients whereas the NW facing node is exposed to high ΔT_{TEG} much less often and therefore harvests most of its energy at lower temperature gradients. We distinguish two operation modes:

net absorption if the overall heat flux is from the ambient towards the ground, and *net emission* when heat from the ground is emitted into the ambient. The NW facing node operates 43.3 % of the time in net absorption, harvesting on average 0.96 mW. The SE facing node harvests on average 2.71 mW in net absorption that occurs during 32.4 % of the experiment duration, and on average 0.32 mW in net emission. The shorter time in net absorption and smaller harvested power in net emission are a result of the experiment duration that was dominated by winter where snow cover insulated the radiator.

Comparison

Due to the numerous design parameters like material, geometry, or thermal contact, to name a few, a detailed comparison is scope of future work. Nevertheless, a comparison with other TEG based systems of with a similar form factor provides a relation to previous work. With an average harvesting power of 1.1 mW, the presented system significantly outperforms the system described in [MSS05]. Their system was slightly larger and generated on average 0.575 mW. Although the average power is identical to what was achieved in [WD12], their system had a larger radiator area. Similarly, the experimentally observed peak harvesting power of 27.2 mW presents the highest value reported in literature for thermoelectric energy harvesting at the ground-to-air boundary with systems that are of comparable size. A systematic comparison to different harvesting modalities, e.g. photovoltaic cells, is highly challenging as numerous factors have to be considered simultaneously. Comparison by area does not suffice, as further aspects including spectral properties, converter architectures, illumination or temperature impact the harvesting efficiency [SC08, SP09]. Developing an appropriate metric and method

that incorporates the necessary factors to enable a detailed comparison between different types of harvesting modalities goes beyond the scope of this thesis.

4.7 Summary

In this chapter, we have introduced an efficient harvesting platform for extracting electrical energy from small bipolar thermal gradients occurring at the ground-to-air boundary. The challenges addressed in the presented work include thermoelectric harvesting from the very low temperature gradients, thermal and electrical matching of thermoelectric generators (TEGs), and efficient rectification of low voltages. The harvesting architecture was modeled end-to-end from ambient conditions including meteorological data to the wireless sensor node application. Furthermore, a novel low-power circuit was designed for rectifying the small bipolar voltages generated by TEGs and incorporated in the overall system model. The model was extensively validated both, component-wise in a controlled lab environment, and overall in a concrete wall scenario. To evaluate its real-world performance, the system was dimensioned and implemented for the use-case of an environment monitoring application. In the long-term case study the harvesting systems was deployed with an wireless monitoring application that senses environmental quantities, participates in a multi-hop network and communicates the acquired information. This experimental evaluation demonstrated self-sustainable operation of the monitoring application with a $550\ \mu\text{W}$ energy footprint. Specifically, the platform harvested up to $27.2\ \text{mW}$ in direct sunlight and $6.3\ \text{mW}$ during night-time, considerably outperforming the state-of-the-art both, in terms of average

and maximum harvested power.

5

Efficient Energy Management for Batteryless Harvesting Systems

While energy harvesting is seen as the key to power cyber-physical systems in a low-cost, long-term, efficient manner, it has generally required large energy storage devices to mitigate the effects of the source's variability. The emerging class of batteryless systems embrace this variability by performing computation in proportion to the energy harvested, thereby minimizing the obtrusive and expensive storage element. As the harvested energy can be too low to satisfy the application's requirements for progress, additional mechanism are needed to operate under such conditions, otherwise the harvested energy is lost. In this chapter, we present an efficient energy management unit (EMU) that decouples the harvesting source and the load, therefore allowing the supply

of generic loads when the average harvesting power is much smaller than required for sustained system operation. By building up charge to a predefined energy level, the EMU can provide short and predictable bursts of energy to a load, even under variable harvesting conditions. Decoupling also allows the load to dynamically adapt its operating point using the proposed dynamic energy burst scaling (DEBS) technique, independent of the harvester's operating point. Experimental data from a solar-powered, long-term autonomous image acquisition application confirms that, regardless of its configuration, the EMU can efficiently supply energy bursts to loads with power demand multiple orders higher than the input power.

5.1 Introduction

Over the past decade, there has been a considerable research effort to reduce the energy consumption of electronic devices. While there has been substantial progress, the lifetime of battery-based devices remains the bottleneck in their development. However, the broader problem of how to supply low power embedded systems with the energy they require in an efficient, low-cost, long-term, scalable, and self-sustainable manner has not yet been adequately solved. Over-provisioning with large energy harvesting and storage elements is either infeasible or unnecessary in many application scenarios such as wearable, autonomous, miniaturized or “smart dust” systems. Fortunately, a purely harvesting-driven system can still meet application requirements in many of these scenarios.

Transiently powered systems are systems that are supplied by volatile energy sources that can, at most, directly

power the system for only a limited amount of time. During this time, the energy harvesting rate can be highly variable, but not necessarily high enough to complete even one atomic task execution, such as performing a sensor reading or transmitting a radio packet. Consequently, such systems need to be able to buffer at least the amount of energy needed to bridge this power deficit and thereby to guarantee the completion of any single task to be executed.

Challenges in Long-Term Batteryless Operation

The design of transient systems must address the challenges listed in the following in order to execute applications in an efficient and reliable manner.

Non-Deterministic Harvesting Environment. Micro-level energy harvesting, which is capable of powering low power wireless sensor networks (WSNs), has received considerable attention in recent years. Interested readers can read [BASM16] for a detailed survey on existing energy harvesting and wireless energy transfer solutions. One important characteristic of all of them, is the intermittent availability of harvested energy. In many cases, harvesting the maximum power from a particular source (e.g. solar, thermal, etc.) requires an impedance matching circuit which dynamically adjusts to changing environmental conditions. Since the system designer has no control over the environmental conditions, no assumptions can be made about the evolution of the source's maximum power point. Unless the complete sensor node and its peripherals are fully scalable in terms of voltage and current, a system which maximizes the input energy has to decouple the source's power point from that of the load.

Variable Load Operating Points and Requirements. Typical low-power cyber-physical systems have components such as microcontrollers, memories and peripherals (e.g. sensors and wireless transceivers). Microcontrollers usually have a wide operating voltage range, but on-chip converters operate most efficiently at lower supply voltages [GPB⁺15]. External peripherals such as sensors and radios can have substantially different voltage requirements, but to minimize cost system designers avoid having multiple voltage domains and simply choose the highest minimum voltage required to supply the entire system. Different tasks, however, can have highly varying energy/voltage requirements and current consumptions since they are highly dependent on the application, peripherals used, etc. In many cases, like ultra low power RFID-scale devices or wearable thermoelectric generator (TEG) harvesting, the operating voltage is so low that high voltage peripherals cannot be used [SWJ⁺13, TSM⁺17]. In order to design a flexible platform that is able to efficiently harvest energy from different sources, it is necessary to decouple the source and load, allowing each to operate at their respective optimal power point.

Minimizing Energy Storage. In many application domains, such as wearable systems, there are stringent form factor restrictions, making storage devices such as batteries and supercapacitors particularly unsuited for long-term use. They are also expensive in terms of cost and area, can have limited charge cycles and high self-discharge rates, impose current peak limits and might not be easily integrated on board [ZGL13]. Consequently, designing a system with an over-dimensioned buffer, if at all possible, invariably leads to

higher losses due to power harvesting costs, additional leakage current, and converter inefficiencies [HSS15]. It is thus indispensable to optimize the energy buffer's size according to application-specific requirements, in order to limit these negative effects.

The aforementioned properties form a veritable challenge from the system design perspective. The energy harvesting has to continuously adapt to the environment conditions to maximize the usable energy. To maximize the application performance, the load needs operate at its optimal operating point. This point can vary significantly due to differing requirements of system components and dynamic power management mechanisms. As a result, the optimal operating points are dynamically changing on both the input and the output of the energy management. Therefore, an energy management architecture should allow independent and dynamic adaptation of these operating points. Furthermore, an application may require a task-dependent minimal energy to progress. Consequently, an energy management architecture should guarantee this energy before the load performs the corresponding atomic step. Designing an energy management architecture that addresses these issues requires innovative methods that combine both hardware and software aspects. In particular, the main challenge is to design a system that can still operate efficiently, have minimized storage and wake-up times given these conditions.

Transient System Configurations

As has been previously discussed, the different properties of transiently powered systems require novel approaches to operate efficiently in such disadvantageous scenarios. In this work, we argue that an additional energy management unit (EMU), as shown in Figure 5.1, is needed to maximize




Figure 5.1: The proposed components (dashed) for efficient and reliable transient system design: an energy management unit (EMU) exposing a control interface for feedback-based dynamic energy burst scaling (DEBS).

the harvested energy and provide the load with the energy guarantees necessary for program progress and at its specified operating point for optimal efficiency. Due to the limited energy intake in transiently powered systems, the unit should self-start requiring as little time and energy as possible. During those periods of limited energy intake, it maximizes the energy build-up by harvesting at the source’s optimal power point. When powering the load with short energy bursts, it should provide a control interface to the load such that the application circuit’s optimal power point can be tracked.

While existing works [LPRR10, Yak11, MMB⁺12, AM16] have looked at low power systems with energy harvesting and storage capabilities, these are expensive in terms of harvesting and storage requirements for long-term, efficient functionality under transient power conditions. State of the art transient system designs [BWM⁺15, JRLR15], connect the energy source directly to the load, without any other intermediaries. However, these works only focus on non-atomic execution of processing tasks, and only work when specific harvesting conditions generate a safe operating voltage on the solar panel.

In this work, we first propose an energy management unit

(EMU) to decouple the operating point of the source from the load. Using an optimally sized buffer capacitor which minimizes the required start-up time and energy from zero, an EMU based system design maintains a low cost, small form factor, high efficiency and virtually unlimited charge cycles. Additionally, a feedback interface is exposed to the load, allowing it to track its optimal power point by applying dynamic energy burst scaling (DEBS). The EMU allows a system to operate predictably and efficiently with limited energy buffering, even under very low power harvesting conditions where the harvested power is significantly lower than the load's minimum power requirement. As an example scenario we consider a solar-powered imaging application. This transient sensor application has the property of guaranteed information and energy availability, since darkness does neither provide energy nor information and light provides both.

Contributions

Specifically, we summarize the contributions made in this chapter as follows:

- We develop the architecture of an energy management unit (EMU) that efficiently converts variable, low power levels to short, high power energy bursts.
- We introduce a feedback-based dynamic energy burst scaling (DEBS) technique to enable a load to track its optimal operating point.
- We optimize the application-specific parameters of the energy management system by means of an accurate model.
- We perform an experimental validation of the high energy efficiency and proportionality of the proposed

transfer scheme in a long-term autonomous image acquisition application.

The experimental results demonstrate that using the highly efficient EMU enables reliable supply of applications where the harvested energy is order of magnitudes lower and highly volatile. Employing the DEBS technique, the total energy consumption of the application can be reduced by 27 % thanks to tracking the optimal power point of its tasks, resulting in a more responsive application.

Outline

The remainder of the chapter is organized as follows: Section 5.2 provides a detailed overview of the state of the art in transient system design. The EMU based energy burst operation and its dynamic adaptation are introduced in Section 5.3. Section 5.4 details the model and architecture of the EMU. The considered application scenario of a long-term vision sensor and its transient system architecture are presented in Section 5.5. The experimental evaluation of the presented mechanisms under different load configurations is provided Section 5.6. We summarize our contributions and findings in Section 5.7.

5.2 Related Work

Cyber-physical system (CPS) supplied from energy harvesting typically included a sufficiently large storage element to mitigate the variability in harvested energy. Due to the prohibitive costs and environmental impact of storing energy, there is a new trend to design systems with minimized storage capacity. Broadly speaking, there are three types of

architectures for transient systems. We cover them in the following sections.

Directly Coupled

When the energy source has an I-V curve compatible with the load, they can be directly connected. The authors of [BWM⁺15] and [JRLR15] have proposed a combined hardware/software approach to perform computation when the source can directly sustain the load during short periods of time. These works use volatile logic that requires state-retention mechanisms. An approach to federating energy proposed in [HSS15] increases the computational ability by using multiple independent capacitors, each dedicated to a specific peripheral. In [KCWP10, WCK⁺14, LC15, WLW⁺16], the authors present storage-less and converter-less harvesting systems in which the load uses frequency scaling to track the maximum power point of the source. While frequency scaling can maximize the energy input in CPU bounded applications, it does not minimize the load's energy consumption and is limited to a narrow active power range. Even though directly-coupled systems avoid converter losses, if the power input is below this restricted active range, the load cannot be powered and the system's efficiency immediately drops to 0%. Unfortunately, this is often the case in typical transiently powered systems. When the energy source and load have incompatible operating points, decoupling them with converters becomes a necessity. In contrast to traditional battery-based systems, these decoupled transient systems have a limited energy buffer between the source and load.

Boost Converter Only

In [DBL⁺15, DLBL⁺16], the authors propose a low-power management system that requires very low input voltage

and current. Using a large buffer capacitor at the converter input, they are able to start the energy conversion at very low input power level. However, both approaches suffer from excessively long cold-start times due to charging a large input capacitance of 140 mF at a constant low input power of 2.5 μ W. As will be explained in Section 5.4.2, our capacitance is chosen to minimize the cold-start energy and start-up time.

Boost/Buck Converter Combination

The authors of [NPK⁺15] also use a boost converter for optimal power point tracking. However, their proposed system utilizes radio frequency (RF) harvesting to accumulate charge in a supercapacitor and then power a camera application with a buck converter. The boost/buck converter topology with an energy buffer also serves as basis for the approach presented in this chapter. While in their work a charge-state model is used to characterize the capacitor's self-discharge rate, energy losses such as impedance matching and converter inefficiencies are neglected. More importantly, the system has a large start-up cost and can only supply the load with bursts of a constant size and voltage. As will be show in the evaluation in Section 5.6, this approach can lead to a substantially higher energy consumption, larger storage elements and longer start-up times.

5.3 Energy Burst-Driven Task Execution

We first overview the energy burst operation principle of the energy management unit (EMU) introduced in this work. Building on the EMU's feedback interface to configure the properties of these bursts, we introduce the dynamic energy burst scaling (DEBS) technique to enable the load to




Figure 5.2: Illustration of the energy management unit (EMU) operation principle: accumulating the variable input energy until the required energy can be provided in an *energy burst* to the load for task execution. Starting from a fully depleted buffer, demands first reaching the load’s minimal supply voltage, leading to a cold-start energy and time overhead (time axis not to scale).

dynamically adjust its operating point. Using this technique allows reducing the application’s energy requirement, thus increasing the efficiency and responsiveness of EMU based transient powered systems.

5.3.1 Predictable Executing using Energy Bursts

Transiently powered systems are designed to operate in limited and volatile energy harvesting scenarios. In order to execute an atomic task, such as reading a sensor value or transmitting a data packet, these systems need to be able to buffer the required energy to guarantee its completion. Consequently, we argue that a novel energy management unit (EMU), as shown in Figure 5.1, is needed to provide energy guarantees in such disadvantageous scenarios. During short periods of limited energy intake, such an unit shall maximize the energy build-up in a small energy buffer by harvesting at the source’s optimal power point, as illustrated in Figure 5.2. The application load is then triggered as soon

as the energy required to run the task to completion has been accumulated. This results in an execution pattern where the energy accumulation from the variable harvesting source is decoupled from providing it as short *energy burst* to the load. This allows separating the concerns of efficient energy harvesting and supplying the application circuit. Due to the limited energy intake in transiently powered systems, the unit should self-start requiring as little time and energy as possible. The details on the architecture and implementation of such an EMU are presented in Section 5.4.

5.3.2 Efficient Task Based Execution

In a typical low power system the load has a varying optimal operating point, as tasks make use of peripherals, such as sensors or communication interfaces, with substantially different voltage requirements. For such scenarios, the proposed EMU shall expose a control interface to enable a load to dynamically adjust the *energy burst* size and voltage to its optimal operating point. Leveraging this feedback capability, we introduce the principle of dynamic energy burst scaling (DEBS): allowing the load to specify the energy required in the next burst and at which voltage it should be provided, the EMU can provide the next *energy burst* at the operating point requested by the load.

To illustrate the operation principle and its benefits, let us consider the example of an application consisting of a two tasks: a sensing tasks τ_1 with a minimal supported voltage of $V_{\text{load},1} = 3\text{ V}$ due to sensor peripheral requirements, and a processing task τ_2 supporting a voltage of $V_{\text{load},2} = 2\text{ V}$. When both tasks should be executed together, this can only be done by generating a static burst at a voltage of 3 V. Consequently, a large energy burst would be required for joint execution due to, a) the joint processing of τ_1 and




Figure 5.3: Splitting the application into tasks τ_1 and τ_2 allows separate execution of them at their respective optimal power point using dynamic energy burst scaling (DEBS). This reduces the energy footprint and increases the responsiveness of the application.

τ_2 , and b) the suboptimal power point for τ_2 . Using the dynamic energy burst scaling (DEBS) technique eliminates the excessive energy consumption due to these two aspects: splitting the execution of τ_1 and τ_2 into separate bursts and specifying their individual requirements using the EMU interface allows executing all tasks at their respective optimal power point. The accompanying reduction in energy requirements allows further optimization of the EMU's internal buffer, as is discussed in Section 5.4.2, thus increasing the overall system efficiency and responsiveness.

Execution of the tasks in separate burst with differing energy and voltage configurations results in a more dynamic energy buffer and task execution behavior, as is shown in Figure 5.3. That employing the DEBS principle leads to significant increases in system efficiency and application execution rate is confirmed by our evaluation in Section 5.6.

5.4 Energy Management Unit

In this section, we describe our model of the proposed energy management unit (EMU) and discuss its architecture shown in Figure 5.4. One of the main goals is to derive equations which can apply to a wide variety of energy sources and loads. The proposed model will then be used to optimize important system parameters, namely the EMU's start-up costs and the load's energy.

5.4.1 Modeling Energy Flow, Buffer and Losses

The amount of energy buffered in the EMU depends on several parameters including the input and load power, and the system's non-idealities. The equation governing the time-dependent energy level in a capacitor is as follows:

$$E'_{\text{buf}}(t) = \frac{d}{dt}E_{\text{buf}}(t) = \eta_{\text{boost}}(V_{\text{in}}(t), I_{\text{in}}(t)) \cdot P_{\text{in}}(t) - P_{\text{load}}(s_i)/\eta_{\text{buck}} - P_{\text{leak}}(t) \quad (5.1)$$

In this equation, the positive term represents the energy intake, while the negative ones represent the energy consumption. The individual components of this model are discussed below.

Input Power

The system has only one power input, $P_{\text{in}}(t)$, supplied by the transducer converting the physical phenomena into electrical energy. This work focuses on the scenario where $P_{\text{in}} < P_{\text{load}}$. In order to maximize the transducer's efficiency, the maximum power point must be tracked to account for variable harvesting conditions, as was illustrated in the characterization results in Chapter 3.

Load Power

In the proposed model, the load has two states s_i : active or inactive. When active, the load is characterized by three quantities: $E_{\text{burst},i}$, $V_{\text{load},i}$, $P_{\text{load},i}$; where $E_{\text{burst},i}$ defines the energy burst size required for one execution of task τ_i , $V_{\text{load},i}$ its supply voltage and $P_{\text{load},i}$ the power consumption during the execution of task τ_i . These parameters have been characterized experimentally. In the inactive state, the load is in deep sleep, consumes very little power, and awaits the trigger from the EMU.

Converter Efficiencies

Since decoupled systems have the source and load operating at different power points, voltage converters are used. This step, while necessary, introduces non-negligible losses, which are represented by boost and buck converter efficiencies $\eta_{\text{boost}}(V, I)$ and η_{buck} . The boost converter's efficiency is particularly sensitive to the variable input voltage and current, meaning it must be parameterized. These efficiencies were characterized experimentally, and interpolation of the values from a look-up table is used in the model. The overall system efficiency of the EMU will be bounded by the product of the boost and buck converter efficiencies. While this depends on both the input and output voltage/current, it goes up to 75 % for the application considered in the experimental evaluation (see Section 5.6).

Additional Energy Losses

Unfortunately, converter inefficiencies are not the only sources of energy losses. The maximum power point tracking (MPPT) unit and the control circuit also consume energy. The consumption of the control circuit I_{ctrl} and buck converter

I_{buck} consist of a constant current and resistive component and hence depend on V_{buf} . For the energy buffer, a capacitor of size C_{buf} , a resistive leakage R_{buf} is assumed. Considering these components, the total system leakage is summarized as:

$$P_{\text{leak}}(t) = V_{\text{buf}}(t) \cdot (I_{\text{ctrl}}(V_{\text{buf}}(t)) + I_{\text{buck}}(V_{\text{buf}}(t))) + V_{\text{buf}}(t)^2 / R_{\text{buf}} \quad (5.2)$$

Equations (5.1) and (5.2) can accurately describe the time evolution of the system's energy levels, as will be shown in the evaluation (Section 5.6). They will be used in the remainder of this section to estimate how different parameters impact the energy losses in the system, to then derive the optimal parameters that minimize these losses.

5.4.2 Reducing Cold-Start Energy and Time

Given the system model presented above, we can start optimizing the cold-start energy and time. By definition this is the fixed cost to turn a transient system on. After a longer period of energy unavailability, the buffer capacitance is fully depleted and first needs to be recharged to the level of V_{load} , as is shown in Figure 5.2. In order to minimize this fixed cold-start cost for a given input power P_{in} , we need to minimize the cold-start time $t_{\text{cold-start}}$ defined as:

$$t_{\text{cold-start}} = \left\{ t \mid V_{\text{buf}}(t) = \sqrt{\frac{2 \int_0^t E'_{\text{buf}}(\tau) d\tau}{C_{\text{buf}}}} = V_{\text{load}} \right\} \quad (5.3)$$

However, the minimum capacitance is limited by the EMU's maximum supported voltage swing, as described by the following equation:

$$C_{\text{min},i} = \frac{2E_{\text{load},i}}{\eta_{\text{buck}}(V_{\text{max}}^2 - V_{\text{load},i}^2)}, \quad (5.4)$$

where $E_{\text{load},i}$ and $V_{\text{load},i}$ are the energy and voltage required to execute a task τ_i of the application, and V_{max} is the EMU's maximum supported buffer voltage. These values must be known at design time, such that the optimal buffer capacitor value can be selected as the highest $C_{\text{min},i}$ among all tasks τ_i , i.e. $C_{\text{opt}} = \max_i\{C_{\text{min},i}\}$. For the implementation the next higher available capacitor size C_{buf} is selected to guarantee task completion.

5.4.3 Minimizing Load Energy

To illustrate the advantages of the EMU's boost/buck architecture compared to a boost-only architecture, let us consider the case of supplying a constant current load consuming I_{load} . The harvesting power P_{in} of a transiently powered system is typically much smaller than the load's power consumption, therefore has a negligible impact on the linear voltage decrease during the time in which the load is supplied with an energy burst. Assuming the load has a maximum supply voltage tolerance from V_{max} down to V_{min} , we have the following power consumption: for the boost-only architecture the average power of a task is $P_A = (V_{\text{min}} + V_{\text{max}})/2 \cdot I_{\text{load}}$, while the buck converter has a constant power of $P_B = (V_{\text{min}} \cdot I_{\text{load}})/\eta_{\text{buck}}$. By comparing these two power consumption values, it directly follows that the buck converter reduces the load's power consumption, if the following condition for the buck converter's efficiency holds:

$$\eta_{\text{buck}} > \frac{2V_{\text{min}}}{V_{\text{min}} + V_{\text{max}}} \quad (5.5)$$

To illustrate this with a numerical example, suppose a load has a voltage tolerance of 3 V to 5 V. This means that a buck converter has a lower power consumption if $\eta_{\text{buck}} > 75\%$. Furthermore, the use of a buck converter offers the possibility




Figure 5.4: The boost/buck architecture of the energy management unit (EMU) with a control circuit exposing the energy burst configuration to the load. An optimally sized capacitor serves as energy buffer for reliable task execution and at the same time reduces the cold-start overhead.

for the load to configure the supply voltage of its optimal power point. When an application consists of multiple tasks τ_i with different voltage requirements, we can use dynamic energy burst scaling (DEBS) to minimize the load's energy requirements.

5.4.4 Energy Management Architecture

The EMU controls the build-up of energy from the source and the energy transfer to the load. Its main components are illustrated in Figure 5.4. The details of their specific implementation are discussed in the following.

Voltage Conversion and Energy Buffering

The harvesting part of the system is based on the commercial bq25505 energy harvesting chip. This chip uses a boost converter to transform the input voltage to a level where the energy can be stored in a storage device. Using its integrated maximum power point tracking (MPPT), the boost converter

adjusts the input impedance such that the power source always operates at its optimal power point to maximize the harvested energy. To provide the required output voltage to the load, the TPS62740 buck converter is directly connected to the energy buffer.

The energy buffer between the input voltage boosting and output voltage regulation guarantees complete separation of the harvesting and load supply unit and therefore allows independent optimization of these parts. As was shown in Section 5.4.2, the storage element minimization is application-specific. The evaluation in Section 5.6 demonstrate with two alternative implementations of the same application how the proposed EMU and DEBS principles support increasing the amount of work completed with the same energy budget.

Control Circuit

The control circuit manages the burst size as well as the output voltage and oversees the energy accumulation in the buffer. For the first, the battery OK signal of the bq25505 is used to trigger the activation of the load, once the capacitor voltage reached a threshold level V_{th} . At this voltage the energy level is reached to provide the requested energy burst to the load. The variable burst size dependent threshold voltage V_{th} is configured using a resistor network. This comparator threshold can be switched digitally from the control circuit by selecting between different resistor networks. Besides very large resistor values, the bq25505 control circuit uses duty-cycling to reduce the energy consumption of the comparator and resistor network. The load supply voltage V_{load} can be controlled directly using the TPS62740 buck converter's digital input exposed through the EMU interface.

Requirements for EMU Operation

Thanks to its inherent decoupling of source and load power points, the minimum requirements for EMU operation are conceptually independent from the load and are only tied to the EMU's circuit implementation. In our case, the first requirement is a minimum input voltage of 330 mV, which is required to turn on a diode in the bq22505 harvester circuit. If this requirement is met, charge is transferred from the harvesting source to a small capacitance (cold-start phase indicated in Figure 5.2). After a certain voltage on this capacitance is reached, the main boost converter is turned on. This transition requires a minimum input current of 60 μ A. This means that as long as the input power is greater than 20 μ W, the EMU is guaranteed to exit the cold-start phase and enter the energy build-up phase. During the build-up phase, the charge on the capacitor will increase as long as $\overline{P}_{\text{in}} > P_{\text{leak,max}}$, where

$$P_{\text{leak,max}} = \{P_{\text{leak}}(t') \mid V_{\text{buf}}(t') = V_{\text{max}}\}. \quad (5.6)$$

After some time, which depends on the input power, enough charge will be accumulated in the capacitor to guarantee completion of the next task and the load is triggered. This makes only the frequency of task activations dependent on the harvesting conditions, but not the task execution itself. Once triggered, the actual task execution is guaranteed, regardless of the harvesting conditions.

5.5 Transient System Architecture

At this point, we have discussed how energy is efficiently harvested and buffered, even in low power harvesting scenarios. The consumption of this energy is, however, a function of

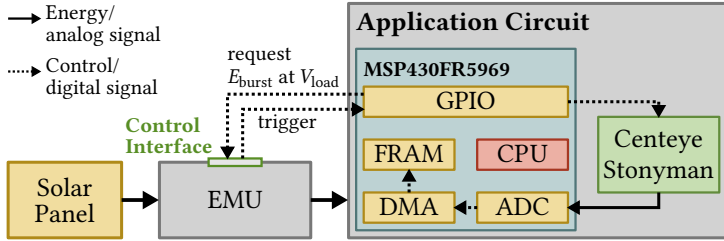


Figure 5.5: The proposed architecture integrates the energy management unit (EMU) with a application circuit consisting of an MSP430FR5969 microcontroller and a Centeye Stonyman image sensor. In the evaluation it is powered from a solar panel.

the load itself. To highlight the flexibility and efficiency of the presented approach, we focus on the sample application of a transient sensor node performing image acquisition and processing. Since the energy management unit (EMU) can work with a wide variety of sources, this section focuses on its integration with the application.

5.5.1 Low-Power Imaging Application (Load)

As an example for a typical sensing load, we use a low power image acquisition application. The architecture of the proposed system is shown in Figure 5.5. The hardware is composed of an MSP430FR5969 microcontroller and a Centeye Stonyman image sensor, which both feature low power consumption and ultra-low power sleep modes. The microcontroller’s IO state lock mechanism and non-volatile ferroelectric RAM (FRAM) are important features to maintain the interface state of the energy manager during deep sleep, and store the task configuration across periods of energy unavailability.

Table 5.1: Energy burst requirements of the considered imaging application when executing its tasks separately, and when executing the entire application in a single burst.

Task/Burst	Voltage (V_{load})	Energy (E_{burst})
Image acquisition task τ_1	3.0 V	215 μJ
Image processing task τ_2	2.0 V	150 μJ
Static application (τ_1 & τ_2)	3.0 V	460 μJ

5.5.2 Energy Burst Configuration

In order to configure the EMU for correct operation, it is important to characterize the application's voltage and energy needs. The considered application consists of two tasks: a) an image acquisition task τ_1 to read the sensor, and b) a basic image processing task τ_2 . The former has an average power consumption of 3.77 mW and minimum voltage requirement of 3 V due to the external camera, while the latter can operate at 2 V and consumes 2.74 mW. Table 5.1 summarizes the energy burst configuration (V_{load} and E_{load}) for the individual execution of the tasks in separate bursts. Further it lists the burst requirements for the baseline configuration without dynamic energy burst scaling (DEBS), where both tasks are executed in a static burst.

As was previously discussed in Section 5.4.2, the EMU's capacitor has to be dimensioned according to the largest task, i.e. the image acquisition. The buffer used in our implementation is a small SMD capacitor of 80 μF . It should be noted that this value is not the optimal theoretical value of $C_{\text{buf}} = 62.5 \mu\text{F}$ calculated using (5.4) to support the largest burst of the baseline configuration, but represents the minimum value required for stable operation of the harvester chip. Therefore, the chosen capacitance minimizes the EMU's cold-start energy and time from a fully depleted buffer capacitor.




Figure 5.6: The software execution flow of an application making use of the dynamic energy burst scheme.

5.5.3 Software Execution Flow

The execution flow of the load's software is shown in Figure 5.6. When the system exits cold-start after a long period of energy unavailability, known as a power-on-reset (POR), the microcontroller performs some basic initialization and immediately enters deep sleep. With a measured power consumption of less than 50 nW, it minimizes losses during the build-up of energy for the next burst. When the next burst is generated, the EMU triggers a control signal to wake up the load. The system then reads the next task configuration and starts its execution after initializing the peripherals needed for that task. At the end of the task, the configuration is updated and the next required burst is configured. Afterwards, the load goes back into deep sleep and waits for the next energy burst to build up. In our evaluation the tasks are repeatedly executed in a static schedule.

5.5.4 Feedback Control for Energy Burst Scaling

As was discussed, there are many application scenarios where the load has a varying optimal power point. This occurs when

tasks use peripherals with substantially different voltage requirements. For such scenarios, the EMU provides a control interface to dynamically adjust the burst size and voltage using the DEBS principle introduced in Section 5.3.2. As this configuration takes only a few microcontroller instructions, its overhead involved is negligible.

Following our image acquisition example, when DEBS is used, the EMU generates two bursts, one for each task. During the first burst, 215 μJ at 3 V are requested. Once enough charge has been built up, the EMU's control circuit configures the buck converter's digital input to set the output to 3 V and triggers the load to execute the image acquisition task τ_1 . Afterwards, the load uses the EMU's interface again and requests the second burst by setting the energy and voltage to 150 μJ and 2 V for task τ_2 . As long as the EMU's buffer has energy, the buck converter will maintain this output voltage until the next burst is generated, the next task executed, and the load requests the next energy burst size and voltage. For the specific properties of the sample application, the two tasks τ_1 and τ_2 are actually executed back-to-back, as the remaining voltage swing after executing task τ_1 at 3 V suffices to start τ_2 at 2 V. Therefore, task τ_2 is triggered immediately after completion of τ_1 and burst reconfiguration.

Without DEBS, the EMU would only be able to generate static bursts at a constant voltage of 3 V. This results in an approach, where one large burst would be used to acquire *and* process the image, similar to the approach proposed in [NPK⁺15]. Such an approach leads to significantly larger burst sizes due to the grouping of tasks with a non-optimal operating point. These two approaches, the baseline using static bursts and the DEBS based execution, will be evaluated experimentally in the following section.

5.6 Evaluation

We evaluate the costs, performance and efficiency of an energy management unit (EMU) powered transient system. For this evaluation we consider the imaging application presented in Section 5.5. To this end, we begin the EMU evaluation by testing its cold-start energy and time, an important result of the optimized buffer size. We then test the performance of the EMU by powering the image acquisition system using two execution profiles: a) *dynamic bursts* using the dynamic energy burst scaling (DEBS) technique, and b) *static bursts* with a fixed, predefined burst configuration. Both execution profiles are tested under constant, as well as variable power input. The experimental results are also compared to a discrete-time simulation of the model presented in Section 5.4 to validate the model accuracy.

5.6.1 Experimental Setup

As sample energy source, the MP3-25 flexible solar panel from PowerFilm with a size of 25 cm² was used. It was exposed only to low illuminance levels of 125 lx to 600 lx, since we focus on low power harvesting scenarios.

Transient Load Configuration

In order to identify the effect of dynamic energy burst scaling (DEBS), two differing configurations are evaluated for prolonged periods of time. The performance of each configuration, under different harvesting conditions, is measured, compared and contrasted.

Static Bursts This configuration serves as baseline for all comparisons. It buffers the energy for the entire application execution in a single burst with constant voltage. This

means that within the same burst the image is acquired and processed.

Dynamic Bursts This configuration, described in Section 5.5.2, uses DEBS for executing the tasks of the imaging application. This means that the first burst is used for image acquisition and the second for processing the acquired data. Each of these bursts is configured to its optimal voltage.

Performance Metrics

In order to compare the performance of different transient configurations, the following metrics are calculated for all experiments:

- $E_{\text{in}} = \int_0^{T_{\text{exp}}} P_{\text{in}}(t) dt$, the total input energy,
- $E_{\text{app},j} = \sum_{i=1}^{N_{\text{tasks}}} \int_{t_{\text{active},i,j}} P_{\text{load}}(t) dt$, the active energy consumed by the j -th iteration of the application,
- $E_{\text{load}} = \sum_j E_{\text{app},j}$, the total energy consumed by the load for all application executions,
- $\eta_{\text{sys}} = E_{\text{load}}/E_{\text{in}}$, the total system efficiency, and
- $U_{\text{exp}} = N_{\text{bursts}}/(N_{\text{tasks}} \cdot T_{\text{exp}})$, the application execution rate.

In the formulas above, $t_{\text{active},i,j}$ denotes execution time of task τ_i in the j -th iteration of the application, N_{tasks} the number of tasks in the application, and N_{bursts} the number of bursts during the total experiment duration T_{exp} . During the experiments, the transient system was instrumented to measure the input and outputs of the EMU, i.e., the input and output currents and voltages, as well as the burst configuration requested by the load.






Figure 5.7: Cold-start time and energy analysis for different input power levels.

5.6.2 Cold-Start Energy and Time Overheads

As was discussed in Section 5.4.2, the energy buffer was optimized to minimize the cold-start energy and time, while still guaranteeing the energy for completion of atomic tasks. To characterize these costs, the capacitor was completely discharged, and the solar panel was exposed to constant illumination level until the cold-start phase ended. The measured cold-start time and energy as a function of the input power are shown in Figure 5.7. The maximum cost, which occurs at the minimum input power of $20 \mu\text{W}$ was 118 s and 2.49 mJ. This was expected since the harvester, by definition, cannot operate efficiently in this region. It should be noted that with an input power of $400 \mu\text{W}$, the cold-start cost go down to 3.9 s and 1.54 mJ. This analysis emphasizes the need for a minimized buffer capacitor, since this also minimizes the cold-start time and energy overhead of the system.

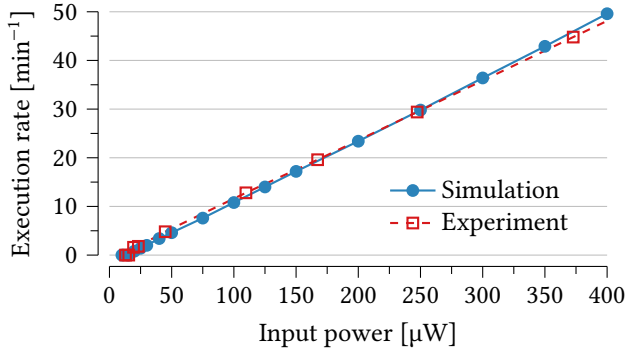


(a) Split task execution, using dynamic energy bursts.

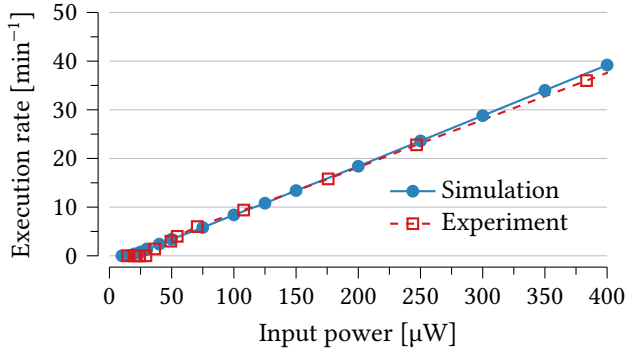


(b) Combined execution using a static burst configuration.

Figure 5.8: Evaluation of the system efficiency η_{sys} under constant input power conditions.



(a) Split task execution, using dynamic energy bursts.



(b) Combined execution using a static burst configuration.

Figure 5.9: Evaluation of the application execution rate U_{exp} under constant input power conditions.

5.6.3 Constant Input Power

In this experiment, the solar panel was exposed to a constant illumination level for 5 min, resulting in a constant power supply to the EMU. This experiment was repeated for different power levels and using dynamic and static burst execution. The system performance was then analyzed using the previously discussed metrics.

The resulting application execution rate U_{exp} and system efficiency η_{sys} are shown for both task execution profiles as a function of the input power in Figures 5.8 and 5.9, respectively. In these plots, we also include the model-based simulation results for the respective scenario. The results show up to 50 and 39 task executions per minute, when using dynamic and static burst execution, respectively. For both profiles, the system efficiency η_{sys} reaches more than 70 % for a wide range of input power levels, with a peak system efficiency of up to 75.1 %. It should be noted that the system model presented in Section 5.4 allows accurate simulation of the number of task executions as well as the overall system efficiency. However, some additional non-linear leakage effects of the boost converter at very low input power of less than 50 μW in combination with high buffer capacitor voltages cannot be captured by the model and result in a too optimistic efficiency for the static burst.

The experimental results show that the execution rate U_{exp} when using DEBS is on average 27 % higher than the static baseline. Further, using dynamic bursts lowers the minimal system operating input power down to 19 μW compared to 36 μW for the static burst execution scheme. Lastly, the system efficiency η_{sys} is increased across the whole input power range for dynamic bursts, with significant improvements for input powers of 200 μW and below.

Table 5.2: System performance analysis for both execution profiles in the real-world experiment with variable input power.

Avg. P_{in}	Metric	Simulation	Experiment	Error
<i>Dynamic energy burst configuration</i>				
92.3 μW	U_{exp}	9.93 min^{-1}	10.33 min^{-1}	-3.9 %
	avg. $E_{\text{app},j}$	368.4 μJ	369.0 μJ	-0.2 %
	E_{load}	54.9 mJ	57.2 mJ	-4.0 %
	η_{sys}	66.11 %	68.82 %	-3.9 %
<i>Static energy burst configuration</i>				
111.9 μW	U_{exp}	9.87 min^{-1}	9.93 min^{-1}	-0.7 %
	avg. $E_{\text{app},j}$	460.8 μJ	459.7 μJ	0.2 %
	E_{load}	68.2 mJ	68.5 mJ	-0.4 %
	η_{sys}	67.76 %	68.01 %	-0.4 %

5.6.4 Variable Input Power

In this experiment the system performance was evaluated in an indoor real-world scenario, again for both task execution profiles. The execution profiles were evaluated in a 15 min experiment that included walking around with the setup in the office hallway partly illuminated by natural and artificial light, walking in a dimly lit basement, and sitting at a well illuminated office desk.

The experimental metrics for dynamic and static burst execution under variable input power conditions are summarized in Table 5.2. First, we note that employing DEBS reduces the average energy per application execution by 19.7 % when compared to the baseline using static bursts. Even though the dynamic bursts experiment had on average a lower input power P_{in} , both, the execution rate U_{exp} and system efficiency η_{sys} , are still higher. This can be explained by the lower energy consumption per task execution due to DEBS’s minimization

of the load's energy consumption. Normalized to the average input power, this results in 22 % more task executions with dynamic bursts compared to static burst execution, which shows the distinct advantage of using DEBS.

Table 5.2 also compares the experimental results to the model simulation that uses the experimental P_{in} data as input. Here, the comparison to experimental values shows that even in a real world scenario with variable input power, the model is able to predict the system performance with maximum error of less than 4 % for both task execution profiles. This fact is also reflected in Figure 5.10: it shows the input power P_{in} , simulated and measured energy level of the buffer capacitor E_{buf} during a 40 s excerpt of the dynamic bursts experiment. Besides a small time drift in the energy accumulation during very low input power, where not all effects can be represented accurately by our model, it tracks the buffer's energy level and bursts with high accuracy. This high accuracy results only in small deviation in the time diagram, despite the accumulation of simulation errors in the time domain.

5.6.5 Result Discussion

The results from the constant power characterization and variable input experiment highlight the four main advantages of our proposed approach. First, thanks to our minimized energy buffer, the cold-start energy and start-up time are minimized: at 400 μW , they were only 1.54 mJ and 3.9 s, respectively. Second, in the very common low power harvesting scenario for transient systems, the EMU decouples the source's and the load's power points. Even though the harvested power never surpassed 400 μW , the EMU reliably provides energy bursts to a 3.83 mW load with a high energy efficiency of 75.1 % when employing DEBS. With a direct coupling approach, it would be simply impossible to power




Figure 5.10: Time domain comparison of the real-world experiment using dynamic energy bursts and the corresponding simulation taking the observed harvesting power P_{in} as input.

the same load. Third, the proposed DEBS technique enables the load to dynamically adjust its operating point and significantly reduce its energy consumption. Finally, the proposed model is able to accurately predict the experimental results. This validates the minimization of our model parameters, namely the minimized energy buffer and cold-start overhead.

5.7 Summary

In this chapter, we presented an energy management unit (EMU) that minimizes the cold-start energy and time for transiently powered systems. By accumulating only minimal amounts of energy in an optimally-sized buffer, the EMU is able to supply generic loads predictably and efficiently, even when harvesting only a small fraction of the load's power requirements. Furthermore, we proposed the dynamic energy burst scaling (DEBS) technique to enable the load to dynamically track its optimal power point. Using an interface exposed by the EMU consisting of only a few digital inputs, the load can dynamically adjust the energy burst size and voltage according to the application's needs, and thus optimize its energy efficiency and performance. The accompanying system model is able to predict the system's performance and energy efficiency within 4.0 % of the experimental values, even under variable power input conditions.

The proposed principles push the limits of energy proportionality by lowering the input power requirements and maintaining a high energy efficiency. Using the EMU and DEBS principles we can build efficient transient systems that operate reliably in highly volatile and unpredictable harvesting environments. In the following chapter these methods serve as basis for implementing an energy-aware communication

scheme for long-term infrastructure-less monitoring.

6

Infrastructure-Less Monitoring with Batteryless Sensors

In this chapter, we focus on infrastructure-less monitoring applications that constitute a highly relevant domain of the emerging Internet of Things (IoT). Operating fully autonomously, batteryless sensor nodes promise truly maintenance-free operation and, therefore, present a compelling solution for long-term monitoring applications. We address the challenge of communicating useful and relevant information in such a scenario, where the senders and receivers operate intermittently due to the variable environment or receiver mobility. Quantifying the importance of sensor information with the notion of data utility, we model this communication scenario and derive an optimized communication scheme. Building on the batteryless energy management and design principles introduced in Chapter 5, we develop a solar pow-

ered Bluetooth Low Energy (BLE) sensor node implementing this communication scheme. Extensive experimental evaluation demonstrates the correctness of the established models and the applicability of the scheme to batteryless sensor nodes with low run-time overhead.

6.1 Introduction

Networked sensing systems are widely used for long-term monitoring of our natural or built environment, personal health, assets, mobility, and many more. While specializing on diverse sensing methodologies and applications, they all accumulate time-dependent sensor information and extract information from received data. Whereas in some scenarios, processing and transmitting the most recent data point is sufficient, the availability of longer time series of data is usually requested and necessary for the end user.

Scenario

We consider long-term monitoring with sensor nodes that target a persistent and maintenance free deployment. These sensors record and locally accumulate environment data, such as temperature, humidity, soil moisture, etc. The sensed data is transmitted to mobile receivers in the vicinity, prioritizing the data most relevant/important to the receivers. To this end, the importance of data is modeled and quantified as a function of its age. The objective is sending a selection of the sensed data that maximizes the usefulness of the data accumulated on the receiver end. Targeting a persistent and maintenance free operation (deploy-and-forget scenario), a suitable communication scheme achieving this goal has to be found. At the same time such a scheme has to consider the

specific requirements of the scenario.

Requirements and Constraints

An *infrastructure-less* communication scheme is required to not rely on deploying and managing any stationary communication or energy supply infrastructure. This facilitates the deployment, operation and long-term maintenance of the sensor network considerably. For *long-term operation* the system needs to operate autonomously for extended periods of up to several decades. *Scalability* of the communication scheme is demanded to allow deployment of many sensors in parallel and to guarantee reception of relevant information with a large number of receivers. These requirements heavily constrain our design and lead to the design choices presented in the following. The mobility of the receiver is a direct consequence of the *infrastructure-less* requirement, since any static data collection device is considered infrastructure.

Energy harvesting is seen as a key enabler for long-term, energy neutral operation [BASM16]. However, the typical combination of harvesting with batteries is limited due to their finite recharge cycles [ZGL13]. To guarantee *long-term operation*, we design our sensor nodes in line with the new class of batteryless devices, which use an energy buffer in form of a capacitor only as in [NPK⁺15, BWD⁺16], or in Chapter 5. Though these systems promise virtually unlimited lifetimes, their operation is highly dependent on the environment. The resulting variability in both timing and energy requires our solution to tolerate and adapt to non-deterministic sensor node execution rates.

To satisfy the *scalability* requirement, we rely on broadcast-based communication, where the sensor nodes advertise data to nearby receivers without employing any request or acknowledgment based scheme. The reason for

this design choice is two fold. Firstly, maintaining a network with batteryless devices is infeasible, as this would require some minimum service level, which cannot be guaranteed in these systems. Consequently, this rules out bidirectional communication schemes in the considered scenario. Secondly, requests and/or acknowledgments do not scale to many receivers with potentially conflicting data requirements. For these reasons, sensor nodes, henceforth referred to as *sensor beacons*, only broadcast a selection of sensed data. The unidirectional communication scheme preserves the privacy of the receivers. Neither a sensor beacon, nor other devices can detect the existence of a receiver.

Challenges

To find an optimal data selection strategy, we introduce the notion of *data utility* that combines the concept of aging of information [ZLT16, HSS17] and the total amount of accumulated data. The challenge in finding the selection strategy that maximizes this *data utility* stems from the uncertainty in two different fronts. Firstly, the selection has to adapt to the non-deterministic activation intervals of batteryless devices that affect the rate at which sensing and transmitting is performed, thus altering the selection of data values providing maximal data utility to receivers. Secondly, mobile receivers are only intermittently within broadcasting range. To maximize the data utility they receive, it is therefore important to incorporate their behavior into the optimization by modeling their listening characteristics. This work is the first to address both of these problems simultaneously and derive the optimal communication scheme.

Contributions

More specifically, we summarize the contributions made in this chapter as follows:

- We propose a receiver-centric communication model for infrastructure-less monitoring in non-deterministic environments with batteryless sensors and stochastic receivers.
- We maximize the receivers' data utility by solving the corresponding convex optimization problem determining the optimized sender-side data selection strategy. This represents the first attempt to optimize the communication policy for infrastructure-less monitoring with batteryless systems.
- Based on a pool of optimized communication policies, we derive a run-time mechanism that adapts to dynamic environmental conditions with negligible overhead.
- We design a harvesting aware batteryless sensor node, implement the proposed communication scheme on top of it, and provide an extensive evaluation that confirms the suitability of the presented models and methods.

The real-world experiments demonstrate the feasibility of the proposed communication and confirm that the run-time mechanisms are implemented with small overhead. With a maximal root mean square error (RMSE) of 0.0157 under controlled power conditions, the experimentally observed data reception probabilities accurately match the theoretical values. The optimized communication scheme is demonstrated to significantly outperform a deterministic baseline scheme in terms of data utility. A final deployment under real-world conditions proves that our scheme successfully adapts the communication policy to the harvesting conditions at run time.

Outline

The remainder of this chapter is organized as follows: Section 6.2 covers related work, before the overall architecture and the communication model are introduced in Section 6.3. In Section 6.4 the optimization problem is formulated and solved, and Section 6.5 presents a number of model generalizations. The batteryless sensor design and implementation of the communication scheme is discussed in Section 6.6. We experimentally evaluate and validate the model in Section 6.7, and summarize our contributions and findings in Section 6.8.

6.2 Related Work

Infrastructure-less sensing has been applied in different scenarios. This includes battery powered wireless beacons for localization [BHE00] and harvesting powered ambient sensor nodes [MMMA05]. Recent batteryless applications include RF sensor tags that deployed cameras [NPK⁺15] or performed ambient sensing [LTLC16]. All of these systems must address issues related to data management, in order to decide what data to communicate. In addition, an appropriate communication scheme must be selected in accordance to the underlying system architecture support. In the following, we address works in these related fields.

Data Utility

To abstract and quantify the usefulness of sensed and transmitted data, several researchers have used the notion of data utility. The authors of [PDMJ10] used data utility to optimize adaptive sensing in wireless sensor networks. Kim et al. [KAS⁺16] applied the notion to data that becomes obsolete after a fixed timeout, and made scheduling decisions

to optimize this utility. Hester et al. [HSS17] introduced a specification language in order to formalize the age dependent utility of data and to perform task scheduling decisions in batteryless sensing. An overview on the use of data utility in information theoretical works [UYE⁺15] focused on offline and online schemes for optimizing the throughput in energy harvesting systems with constrained energy buffers. A more recent work in this domain [FY18] investigated optimal sampling and delay minimal communication strategies for energy harvesting system in single user and broadcast scenarios.

Following the same idea of age of information, we rely on the notion of data utility to derive an optimization formulation and determine the optimal communication strategy for infrastructure-less sensing scenarios.

Data Compression

Data utility can be maximized using compression, as the sensed data is typically correlated in time and space. The compression allows for significant communication energy savings with reasonable processing overhead [SM06]. The methods used in wireless sensor network can be broadly grouped as follows [RBD13]: entropy coding based compression [CK18], prediction based compression [BMG18], transformation based compression [Mer05] or applying compressed sensing [CW08].

The communication scheme presented in the following leverages the low complexity Haar wavelet transformation for lossy compression of long-term historical values.

Infrastructure-Less Communication

Specialized communication schemes are required to disseminate data without relying on any infrastructure like wires

or base stations. In the area of delay tolerant networks, researchers addressed the mobility on the sender and/or receiver side [LYQ11, YZLM14]. The various data dissemination schemes proposed to handle the intermittent connectivity are based on flooding or routing and make use of known or learned mobility patterns and neighbor discovery [SPR05, WTZE09, RN17]. Similarly, energy harvesting sensor networks employed flooding [JOW⁺02, MMA05] or routing based schemes [MKK⁺15], relying on large batteries to mitigate the non-determinism in energy availability. Typically, these systems adapted the communication and sensing rate jointly to the harvesting conditions to sustain system operation [GWZ13, SBC⁺13].

These techniques are not directly applicable to batteryless systems, as they rely on periodic discovery and updates with neighboring nodes for dynamic forwarding and routing. We therefore employ a broadcasting communication scheme.

Batteryless System Support

The emerging class of batteryless systems avoids bulky batteries as energy supply and solely relies on energy harvesting and small energy buffers such as supercapacitors. Numerous proposals focusing on node architecture challenges have been published. They addressed the non-deterministic system operation by introducing state retention with periodic [RSF11] or event-driven state saving mechanisms [JRLR15]. Others used an energy-driven approach by providing the energy in small, but guaranteed bursts and extending the hardware for more flexible energy management such as [NPK⁺15, CRL18] and the concepts introduced in Chapter 5. Furthermore, the problem of persistent timekeeping has been addressed [HTR⁺16] and higher level abstraction for operating system integration and data management have recently been proposed [BM17,




Figure 6.1: The considered scenario consists of a batteryless sensor node monitoring the environment and broadcasting (RF) data packets of N sensor values to smartphones with listening intervals described by distribution $\tilde{\mathbf{R}}$. The application specification $\tilde{\mathbf{Q}}^{\text{spec}}$ describes the relative importance of individual data units.

HSS17, BMAS19].

The batteryless sensor developed as part of the presented work builds on the energy burst principle introduced in Chapter 5. Furthermore, a backup power domain is incorporated so that sensor nodes do not lose the notion of time over long periods of energy unavailability, for example during nighttime.

6.3 System Overview

First, we introduce the infrastructure-less communication scenario more formally in this section. We state the models required for formulating the optimal communication problem, followed by an architectural overview in which we introduce the different components of the proposed communication scheme.

6.3.1 Communication Scenario

We focus on a long-term infrastructure-less ambient monitoring scenario, as illustrated in Figure 6.1: a sensor node

accumulates ambient sensor data and buffers them as time-stamped *data units* in non-volatile memory. The sensor node broadcasts a selection of N data units per data packet, which can be received by devices within its range. There may be several receivers active simultaneously collecting this data. We consider smartphones as receiver devices, as their ubiquity suits the infrastructure-less sensing scenario very well.

The scenario supposes that a receiver values the availability of data depending on their *age*. To this end, a specification vector $\tilde{\mathbf{Q}}^{\text{spec}}$ is introduced that enables an application designer to specify the relative importance of a data unit as a function of its age. Considering air quality monitoring as an example, this could be reflected by a specification with an importance that decreases with age, as represented by the *logging* specification shown in Figure 6.2. For applications that provide feedback or control based on the received data, typically only values up to a maximum age are considered useful, as is reflected by the step function of the *feedback* specification. For long-term monitoring where the availability of a long history is of higher importance than detailed coverage, a specification as presented by *monitoring* in Figure 6.2 could be used: data units with a large interval have higher importance, while the importance decreases for intermediate, finer granular data.

Similarly, the mobility of receivers is specified by a listening interval distribution $\tilde{\mathbf{R}}$. It characterizes the duration a specific receiver is within the range of the sensor and able to receive data. Examples for this specification are shown in Figure 6.3. An *office* scenario reflecting the mobility of an office worker, a *meeting room* scenario in which people are present for one or two hour meetings, or a *kitchen* scenario in which people are either briefly dropping to get a coffee, or having an extended lunch break.

The system specification $(\tilde{\mathbf{Q}}^{\text{spec}}, \tilde{\mathbf{R}}, N)$ allows for a quan-




Figure 6.2: Examples of the relative importance specification \tilde{Q}^{spec} for a few reference scenarios with differing requirements.




Figure 6.3: Examples of receiver characteristics \tilde{R} for three different scenarios.

titative description of a scenario and enables us to derive an optimized communication strategy for the sensor node. The exact objective of the corresponding optimization problem, i.e., the exact notion of *data utility* used in this work, is defined in Section 6.4.1.

6.3.2 Notation and Communication Model

We introduce a model that abstracts the considered infrastructure-less monitoring scenario. The model consists of two parts: the sensor node, also referred to as *sender*, that transmits a selection of the sensed data, and a (potentially large set of) device(s) referred to as *receiver(s)* that collect the broadcast data.

Notation

Throughout this chapter bold symbols are used to denote vectors and matrices and subscripts denote individual elements of the vector or matrix. For instance, if \mathbf{X} is a vector, X_i denotes its i^{th} element. In case \mathbf{X} is a matrix, \mathbf{X}_i denotes its i^{th} row, and $X_{i,j}$ denotes the element in its i^{th} row and j^{th} column. For a discrete random variable \mathcal{Y} , \mathbf{Y} denotes the vector of probabilities, whose i^{th} element is the probability that \mathcal{Y} is equal to i ; i.e. $Y_i = \Pr(\mathcal{Y} = i)$. Y^* denotes the largest index i for which $Y_i \neq 0$. All time dependent model components are specified in terms of *unit time* δ_0 that defines the time granularity.

Sender

The sender performs sense-process-transmit operations based on the principle of energy bursts. Its operation is characterized by the following parameters:




Figure 6.4: The overall architecture consists of an *online* harvesting aware data transmission scheme. The key component is a pool of communication policies that are optimized *offline* based on a model for the infrastructure-less communication scenario.

- δ Integer specifying the time between sender activations, defined as a multiple of the base time unit δ_0 . For the presentation of the model and the optimization we assume that $\delta = 1$, i.e., the activation interval corresponds to the base time unit δ_0 . This can be generalized to arbitrary δ as discussed in Section 6.5.1.
- N Size of a data packet in number of data units.
- $E^{\text{activation}}$ Energy required per sensor node activation. This is the sum of the energy to sense a new data unit E^{sense} , buffer management and data selection for transmit E^{process} , and sending one data packet of N data units E^{send} .
- pwr^{in} Input power available to the sensor node. This is the harvesting power actually available after accounting for transducer, conversion, energy management, and storage inefficiencies.

Given these parameters, the following condition on the average input power $\widehat{\text{pwr}}^{\text{in}}$ must be satisfied for sustainable sensor node operation:

$$\widehat{\text{pwr}}^{\text{in}} \geq \frac{E^{\text{activation}}}{\delta \cdot \delta_0} \quad (6.1)$$

The interval δ has to be adapted accordingly by a run-time mechanism to satisfy the above condition.

The selection of the N data units is determined by the probability vector \mathbf{P} . The individual P_i specify the probability of including data that have been sensed at time $i \cdot \delta$ before the current time.

Receiver

The receivers arrive at an unknown point in time and are active for a limited time only. As introduced above, the

individual probability of occurrence of the listing intervals is described by vector $\tilde{\mathbf{R}}$. Furthermore, the specification $\tilde{\mathbf{Q}}^{\text{spec}}$ states the relative importance of data units for a receiver. Depending on the activation interval described by δ , the specifications $\tilde{\mathbf{Q}}^{\text{spec}}$ and $\tilde{\mathbf{R}}$ are rescaled to \mathbf{Q}^{spec} and \mathbf{R} . To understand why the rescaling is necessary, consider the following example: if a sensor is activated every minute, the entries defined for $[1, 2, \dots, 59]$ seconds of a $\tilde{\mathbf{Q}}^{\text{spec}}$ with second granularity are obsolete, since the corresponding data is never sensed. Therefore, we need to adapt $\tilde{\mathbf{Q}}^{\text{spec}}$ and $\tilde{\mathbf{R}}$ based on activation interval δ . The details of this rescaling step are explained in Section 6.5.1. For the assumption of $\delta = 1$, the probability vectors \mathbf{R} and $\tilde{\mathbf{R}}$, and \mathbf{Q}^{spec} and $\tilde{\mathbf{Q}}^{\text{spec}}$ are identical. The receiver side reception probability is encapsulated in \mathbf{Q} , with its elements Q_i representing the probability of receiving a data unit that is $i \cdot \delta$ time units old. They depend on both, the receiver characteristics \mathbf{R} and selection probabilities \mathbf{P} .

In the following, the sensor node architecture and its operation is detailed, followed by an overview of the communication scheme optimization that is discussed in depth in Section 6.4.

6.3.3 Sensor Node Architecture

The detailed operations of the sensor node performed as part of an activation are illustrated in the *online* part shown on the right in Figure 6.4. The node starts with reading the ambient *sensors* and *timer* and then stores these values as one data unit in a non-volatile *buffer*. In addition, aggregate values like a compressed history are calculated in the *update aggregates* component. At the end of an activation, the node broadcasts (RF) a selection of N data units. The data units to transmit are selected in the *packet selector* unit according to the optimal selection policy ρ .

The optimal policy is dependent on the harvesting condition pwr^{in} that dictates the time interval between activations (6.1). Therefore, a *policy pool* is stored on the sensor node, containing optimal policies for different activation intervals δ . At run time, a *duty-cycle adaptation* calculates the activation interval that allows sustainable operation of the sensor and activates the corresponding optimal policy ρ . The derivation of the optimal policy for individual intervals is done offline based on the procedure outlined below.

6.3.4 Determination of the Policy Pool

The optimal communication policy depends on the system specifications $(\tilde{\mathbf{Q}}^{\text{spec}}, \tilde{\mathbf{R}}, N)$ and the current input power pwr^{in} . An energy model of the sensor node is used to calculate the activation interval δ that guarantees long-term sustainable operation for a constant input power pwr^{in} .

The scaled specifications \mathbf{Q}^{spec} and \mathbf{R} are passed together with the data packet size N to the model based *optimization*. This block solves the optimization formulation detailed in Section 6.4 and returns the resulting vector \mathbf{P} specifying the optimal probabilities of selecting data units of a certain age. The probabilities are bin-packed to get individual probability mass functions for each of the N data slots of a packet. The resulting probability mass functions are summarized as the optimal selection policy ρ and used at run time to select the N data units to be transmitted.

The optimization flow of determining δ , rescaling $\tilde{\mathbf{Q}}^{\text{spec}}$ and $\tilde{\mathbf{R}}$, finding the optimal probabilities \mathbf{P} , and determining ρ by bin-packing \mathbf{P} , is repeated for several discrete power levels. The corresponding set of optimal communication policies for different activation intervals δ define the policy pool $\{(\delta, \rho)\}$ that is deployed on the sensor node.

6.4 Optimal Communication Policy

Based on the specifications and the communication model introduced before, we derive the optimal data selection policy. First, we formalize the notion of *data utility* that is serving as optimization objective. We then analytically derive the probability Q_i of receiving data with an age of i time units. Subsequently, we present the convex optimization problem to determine the optimal sender side data selection probability. Finally, the online data selection algorithm for implementing the optimal communication policy is introduced.

6.4.1 Data Utility

The data utility accumulated for given values of \mathbf{Q} and \mathbf{Q}^{spec} is given by the following equation:

$$U(\mathbf{Q}, \mathbf{Q}^{\text{spec}}) = \min_{i \in \{0, \dots, Q^{\text{spec}*}\}} \{Q_i / Q_i^{\text{spec}}\} \quad (6.2)$$

This utility formulation takes into account the relative importance \mathbf{Q}^{spec} of data of a given age. Higher values of Q_i / Q_i^{spec} mean that important data is received with higher probability. Note that maximizing the above utility $U(\mathbf{Q}, \mathbf{Q}^{\text{spec}})$ for a given \mathbf{Q}^{spec} is identical to the following formulation:

$$\text{maximize } \alpha \quad \text{s.t. } Q_i \geq \alpha \cdot Q_i^{\text{spec}} \quad \forall i \in \{0, 1, \dots, Q^{\text{spec}*}\} \quad (6.3)$$

i.e. determining a \mathbf{Q} such that all its elements are no less than a maximally scaled version of \mathbf{Q}^{spec} .

6.4.2 Probability of Data Reception

To derive the data reception probability Q_i we consider an introductory example. Let us examine a receiver that is within the range of a sensor node for 3 time units, i.e., at times

0, 1, 2. After time 2 the receiver moves out of range and cannot receive further data packets. Furthermore, we assume the sensor transmits one data packet every interval. To find the probability of having received data of age 5 at least once when moving out of range, there are three cases to consider:

- a) Data of age 3 is sent at time 0. In this case, the data ages for 2 time units on the receiver side and therefore has an age of 5 when the receiver moves out of range.
- b) Data of age 4 is sent at time 1 and ages for 1 time unit on the receiver side.
- c) Data of age 5 is sent at time 2.

Each of these cases occurs with a specific probability that depends on the selection probabilities P . The data reception probability Q_i of i time units old data on the receiver side is influenced by both, the receivers listening interval distribution R and the data selection probabilities P . For the analytical derivation of Q_i presented here, the communication link is assumed lossless for simplicity reasons. An extension of the model to incorporate a time-invariant packet reception probability for lossy communication channels is presented in Section 6.5.4.

To derive Q_i we look at the $i + 1$ scenarios in which an i time units old data unit is not received. These scenarios are represented in Figure 6.5: the enumerated gray nodes refer to the active time of the receivers, while the double circled red nodes indicate the individual scenarios in which the i time units old data is not received. While active, receivers transitions horizontally from the circled initial node to any of the gray nodes enumerated with the receivers' active time. The double circled red nodes are reached by vertical transitions referring to cases in which a receiver becomes inactive after the number of intervals indicated by the label of its source node. The probability of a scenario s_i is calculated by




Figure 6.5: Enumeration of the scenarios where i time unit old data is not received (s_0, \dots, s_i).

multiplying the probabilities indicated on the edges leading to that scenario. In the first scenario s_0 , a receiver is active for exactly one time unit and the sender has not selected the i^{th} data for transmission. The probability of this scenario is $(1 - P_i) \cdot R_0$. Scenario s_1 encapsulates the situation where a receiver is active for exactly two time units and the sender does not transmit the i^{th} data unit at time 0 and the $(i-1)^{\text{th}}$ data unit at time 1. The last scenario s_i refers to a receiver active for *at least* i time units and the sender does not send i^{th} data at time 0, $(i-1)^{\text{th}}$ data at time 1, and so on. Summarizing all scenarios of not receiving i time units old data, the probability Q_i is calculated as:

$$\begin{aligned}
 Q_i &= 1 - \underbrace{\sum_{j=i}^{j \leq R^*} R_j \cdot \prod_{j=0}^{j \leq i} (1 - P_j)}_{\text{scenario } s_i} \\
 &\quad - \underbrace{\sum_{m=1}^{m \leq i} \left(R_{i-m} \cdot \prod_{j=m}^{j \leq i} (1 - P_j) \right)}_{\text{scenarios } s_0, \dots, s_{i-1}}
 \end{aligned} \tag{6.4}$$

The first sum in the formula computes the probability of scenario s_i , while the second sum covers the remaining i scenarios (s_0, \dots, s_{i-1}) by summing up the respective proba-

bilities of occurrence for each scenario. The probabilities of all scenarios can be accumulated because all scenarios are disjoint. We also use the following equivalent formulation for conciseness:

$$Q_i = 1 - \sum_{j=0}^{j \leq R^*} \left(R_j \cdot \prod_{k \in S(i,j)} (1 - P_k) \right), \quad (6.5)$$

where the sets of indexes $S(i, j)$ are defined as:

$$S(i, j) = \begin{cases} \{i - j, \dots, i\} & j < i \\ \{0, \dots, i\} & \text{otherwise} \end{cases} \quad (6.6)$$

6.4.3 Optimization of Data Reception

The optimal data selection probability vector \mathbf{P} is derived using an optimization formulation based on the specified reception probability \mathbf{Q}^{spec} , receiver listening interval distribution \mathbf{R} , and packet size N . For simplicity, it is still assumed that $\delta = 1$. A generalization to arbitrary activation intervals δ is discussed in Section 6.5.1.

The optimization problem is stated as follows:

$$\text{maximize: } \min_{0 \leq i \leq Q^{\text{spec}*}} \{Q_i / Q_i^{\text{spec}}\} \quad (6.7)$$

$$\text{subject to: } 0 \leq P_i \leq 1 \quad \forall i \in \{0, \dots, P^*\} \quad (6.8)$$

$$\sum_{0 \leq i \leq P^*} P_i \leq N - \epsilon \quad (6.9)$$

The objective function (6.7) maximizes the minimum ratio between the elements of \mathbf{Q} and the corresponding element of \mathbf{Q}^{spec} . Therefore, the objective maximizes the data utility $U(\mathbf{Q}, \mathbf{Q}^{\text{spec}})$ specified in (6.2). The constraints specified in (6.8) ensure that all elements of \mathbf{P} are valid probability values. Constraint (6.9) ensures that the number of data units selected does not exceed the data packet size N . The parameter ϵ is used

Algorithm 1: First Fit Decreasing

```

1 Function FirstFit( $\mathbf{P}, N$ ):
2    $\rho_{i,j} \leftarrow 0 \quad \forall i \in \{1, \dots, N\}, j \in \{0, \dots, P^*\}$ 
3   for  $(i, P_i) \mid P_i \in \mathbf{P};$  in descending order of  $P_i$  do
4     for  $j = 1$  to  $N$  do
5       if  $\sum \rho_j \leq 1 - P_i$  then
6          $\rho_{j,i} \leftarrow P_i$ ; mark  $P_i$  as assigned
7         break
8     if  $P_i$  is not marked as assigned then
9       return failure
10  return  $\rho$ 

```

to strengthen the requirement of at most N data units per data packet. This additional slack is used in the bin packing stage that is described in the next section, see also Figure 6.4.

The optimization formulation specified in (6.7) to (6.9) is proven to be convex (see proof in Section 6.A.1). Given the convexity of the problem, it can be solved efficiently and optimally by any of the available convex optimization solvers [BV04]. A feasible solution is returned as long as $N - \epsilon \geq 0$.

6.4.4 Optimal Data Selection Policy

The optimal selection probability \mathbf{P} of data units cannot directly be used to decide whether a data unit is transmitted in a given interval, as this would lead to a variable data packet size. Therefore, a final partitioning step is added to distribute these probabilities to the N slots of a fixed size data packet. To guarantee that no duplicates are selected for transmission, First Fit bin packing (Algorithm 1) is used to determine individual probability mass functions for each of the N data slots of a data packet. The parameter ϵ in (6.9) is

used to account for partitioning loss, so that no more than the available N data slots are used. The probability mass functions for each of the N slots of a data packet are summarized in a data selection policy matrix ρ , where rows ρ_i refer to the probability mass function for selecting the data unit of the i^{th} data slot.

The exact procedure for obtaining an optimal data selection policy ρ is described in Algorithm 2. The slack variable ϵ is initialized to zero (Line 1). This is followed by finding the optimal selection probability vector \mathbf{P} using the optimization formulation (Line 2). The optimizer returns \mathbf{P} such that $\sum_i P_i \leq N - \epsilon$. Then the algorithm attempts partitioning \mathbf{P} into N unit sized bins using First Fit Decreasing bin packing (Algorithm 1). In case the partitioning succeeds, a matrix ρ with its N rows being the required N partitions is obtained. The algorithm normalizes each row of ρ to receive valid probability mass functions (Line 6) and returns the resulting ρ . If bin packing is not successful, i.e., more than N bins are required, then ϵ is incremented by an arbitrary, predefined step of 0.01 and the procedure is repeated.

At run time, the sender makes the i^{th} selection decision by randomly sampling from the resulting ρ_i (see also Figure 6.4). Since, each element of \mathbf{P} exists in exactly one probability mass function, it is guaranteed that each data unit is selected at most once per data packet. Therefore, the calculation of Q_i (the probability of receiving data of age $i \cdot \delta$) in (6.5) holds when packet selection is done using ρ as explained here.

After discussing a number of generalizations in the next section, the implementation of the run-time mechanisms on a real-world batteryless sensor node is discussed in Section 6.6.

Algorithm 2: Optimal Data Selection Policy

Input: $\mathbf{Q}^{\text{spec}}, \mathbf{R}, N$
Output: Matrix ρ with N rows of probability mass functions

```

1 for  $\epsilon = \{0, 0.01, \dots, N\}$  do
2    $\mathbf{P} \leftarrow$  Solution of the optimization problem given
   in (6.7) to (6.9)
3    $X \leftarrow \text{FirstFit}(\mathbf{P}, N)$ 
4   if  $X \neq \text{failure}$  then
5      $\rho \leftarrow X$ 
6      $\rho_{i,j} \leftarrow \frac{\rho_{i,j}}{\sum \rho_i} \quad \forall i \in \{1, \dots, N\}, j \in \{0, \dots, P^*\}$ 
7     return  $\rho$ 

```

6.5 Model Generalizations

In the following we generalize the model to allow for different time units δ , add support for aggregated sensor values, introduce run-time harvesting awareness, and to incorporate communication over a lossy channel.

6.5.1 Relaxing Activation Interval

The parameter δ characterizes the interval between consecutive activations of the sender. While $\delta = 1$ was assumed in the presentation of the model, in reality it needs to be adapted to the harvested input power pwr^{in} to guarantee sustainable operation. For this reason, a mechanism for scaling the specification of the receiver behavior $\tilde{\mathbf{R}}$ and the requirement \mathbf{Q}^{spec} to arbitrary δ is needed.

Remember that $\tilde{\mathbf{Q}}^{\text{spec}}$ and $\tilde{\mathbf{R}}$ are defined for the lowest time granularity δ_0 . As the interval δ is defined to be an integer multiple of δ_0 , the scaled model inputs \mathbf{Q}^{spec} and \mathbf{R} can be received by down-sampling the specifications.

Several aggregation options like mean, min or max could be used for \tilde{Q}^{spec} . In this work, the probability of receiving a packet at any point in time within the down-sampling interval is used, corresponding to the following equation:

$$Q_i^{\text{spec}} = 1 - \prod_{j=i\cdot\delta}^{j<(i+1)\cdot\delta} (1 - \tilde{Q}_j^{\text{spec}}) \quad (6.10)$$

This avoids relaxing of the specification through aggregation and smooths hard transitions for large intervals.

To receive a valid \mathbf{R} that represents a listening interval probability mass function, the intermediate probabilities of $\tilde{\mathbf{R}}$ are accumulated for each aggregation interval:

$$R_i = \sum_{j=i\cdot\delta}^{j<(i+1)\cdot\delta} \tilde{R}_j \quad (6.11)$$

Therefore, when specifying Q^{spec} or \mathbf{R} in the model, it is implicitly assumed that these values are determined using equations (6.10) and (6.11) for the contextual value of the activation interval δ .

6.5.2 Aggregated Sensor Values

While the introduction of the model focused on transmitting aging sensor values, the model can be extended to incorporate aggregate values. More specifically, we consider lossy compression of the long-term history using the Haar wavelet transformation [Mer05]. Among the resulting coefficients the most significant ones are transmitted, with a decreasing probability specification.

To include P^{agg^*} aggregate values, the selection probability vector \mathbf{P}^{agg} is introduced, where its individual components P_k^{agg} define the probability of including the k^{th} aggregate value in a data packet. The components of the corresponding reception probabilities Q^{agg} are calculated analogously to the

probability of receiving aging values in (6.5). In contrast to aging values, the calculation is time-invariant:

$$Q_k^{\text{agg}} = 1 - \sum_{j=0}^{j \leq R^*} R_j \cdot (1 - P_k^{\text{agg}})^{j+1} \quad (6.12)$$

By augmenting the vectors \mathbf{P} , \mathbf{Q} and \mathbf{Q}^{spec} with the corresponding variables and specifications for the aggregated values \mathbf{P}^{agg} , \mathbf{Q}^{agg} , and $\mathbf{Q}^{\text{spec,agg}}$, the original optimization formulation given in (6.7) to (6.9) is extended to support aggregate values. Note that in the calculation of Q_i for aging values (6.5), only elements of \mathbf{P} which are not in \mathbf{P}^{agg} are used. This extension does not impact the convexity of the optimization formulation (see proof in Section 6.A.2). Optimal policy generation using Algorithm 2 is therefore guaranteed. Neither the algorithms, nor the run-time selection mechanisms need to be updated to support aggregate values. The only addition needed is the actual calculation of the aggregate values at run time.

6.5.3 Energy Harvesting Awareness

Inherent changes in the environment lead to a time varying input power pwr^{in} . Changes in pwr^{in} demand an update of δ to satisfy condition (6.1) for sustainable operation and consequently a model change. To support energy harvesting aware data selection, the optimal selection policy calculation is repeated for a set input power levels in the offline optimization: using the energy model of the sensor node, the minimum δ for long-term sustainable operation under the assumption of a constant power level is computed based on (6.1). With the specifications $\tilde{\mathbf{Q}}^{\text{spec}}$ and $\tilde{\mathbf{R}}$ scaled accordingly, the schedule is optimized using Algorithm 2. The generated optimal selection policy $\boldsymbol{\rho}$ and the corresponding δ for the considered power levels are summarized in a policy pool and deployed on the

sensor node.

For selecting the optimal policy at run time, the sensor node deploys a duty-cycle adaptation mechanism that updates δ . The adaptation and energy estimation procedures used for this are beyond the scope of this work, we refer the reader to existing run-time adaptation methods [KHZS07, MTBB10, BSBT14b, ABD⁺19]. To adapt the communication policy accordingly, the optimal selection policy ρ with the activation interval closest to the new duty-cycle is loaded from the policy pool and forwarded to the run-time packet selection.

6.5.4 Lossy Communication Channel

We introduce a time-invariant packet reception rate $PRR \in [0, 1]$, specifying the probability that a packet is received and decoded correctly, to characterize a lossy communication channel. This metric is integrated in the overall model by replacing the term $(1 - P_k)$ with $(1 - P_k \cdot PRR)$ in (6.5), leading to the lossy data reception probability of

$$Q_i = 1 - \sum_{j=0}^{j \leq P^*} \left(R_j \cdot \prod_{k \in S(i,j)} (1 - P_k \cdot PRR) \right). \quad (6.13)$$

With PRR being a constant scaling parameter in the closed unit interval, its addition does not affect the convexity of the optimization formulation specified in (6.7) to (6.9).

6.6 Sensor Node Implementation

For testing the proposed communication scheme in a real-world sensing scenario, we designed a custom batteryless




Figure 6.6: The transient Bluetooth Low Energy (BLE) sensor node is based on the CC2650 system on a chip (SoC). The SoC interfaces a real-time clock (RTC) with separate backup power domain and a ferroelectric RAM (FRAM) to provide timekeeping capabilities and data retention across power failures. For batteryless operation, the node relies on the energy management unit (EMU) and energy burst principle introduced in Chapter 5 and a solar panel as energy source.

Bluetooth Low Energy (BLE) sensor node. Any BLE enabled Android smartphone can receive and decode the sent data using our receiver application.

6.6.1 Transient Sensor Node

Being a transient sensor node, the energy buffer is dimensioned to only support the minimum atomic operation that guarantees progress, therefore belonging to the most restricted share of batteryless systems. The execution is fully dictated by the environment, as the device accumulates ambient energy in a small buffer until enough energy is available for the next step of the application.

Hardware Platform

The transient BLE sensor node is based on the Texas Instruments CC2650 system on a chip (SoC). Its architecture is detailed in the diagram in Figure 6.6. For batteryless




Figure 6.7: The custom designed transient sensor node used to implement and evaluate the communication scheme has a small size of only 35 mm × 53 mm, about the size of the match-box shown in the background. It is powered from the solar panel mounted on the back of the node.

operation, the energy management circuitry presented in Chapter 5 is integrated with a small energy buffer of $C_{\text{buf}} = 200 \mu\text{F}$ and configured with a constant energy burst size. The sensor node is deployed with an AM-5412 solar panel of 50 mm × 33 mm size covering the back of the node to harvest energy from ambient light. A backup power domain for the external AM0815 real-time clock (RTC) enables persistent timekeeping across several hours of energy unavailability using a backup capacitor of $C_{\text{bak}} = 320 \mu\text{F}$. The buffer of this backup domain is recharged as part of an activation of the sensor node. An external FM25V10 ferroelectric RAM (FRAM) memory is included for energy efficient storing of the sensor data history, aggregate values, and system state. For sensing the ambient temperature and humidity an SHT31 sensor is integrated on the node platform. The transient sensor node was implemented as a custom printed circuit board (PCB) that is shown in Figure 6.7.

Software

The application logic is implemented on top of the Contiki operating system [DGV04] that comes with the required radio drivers to send standardized BLE data packets. Driver support for interfacing the energy management unit (EMU), sensor, timer and FRAM was added to the operating system and the start-up sequence was optimized for fast and energy efficient start-up. The following sequence of operations are performed by the sensor node during one activation (see also Figure 6.4):

1. Restore the system and data buffer states from the FRAM,
2. Increase the system voltage to 3.3 V and start recharging the timer backup domain's capacitor C_{bak} ,
3. Readout the current time and sensor data, and store them as a new data unit in the FRAM data history,
4. Once a day, compress the accumulated 6 min aggregates using the Haar wavelet transform and store the 10 most relevant coefficients as aggregate values,
5. Estimate the time-based power level to load the corresponding optimal policy ρ from the policy pool,
6. Sample probabilistically according to ρ to get the selection of data units to include in the next data packet,
7. Load the selected data units from the FRAM, assemble them in a BLE packet, and broadcast it on all BLE advertisement channels (CH37-CH39),
8. Back up data buffer and system state and enter deep sleep at the minimal system voltage of 2.3 V.

In cases of extended periods of energy unavailability the RTC backup domain can run out of energy and the timer is reset on the next activation. In such a case the data buffer needs to be cleared due to the loss of the time reference of the buffered data units.

6.6.2 Smartphone Receivers

Commodity smartphones, Samsung Galaxy A3 (2016), were used as receiver devices for data collection. Our application scans for BLE advertisement packets in the background and logs timestamped packet data to internal memory. The collected traces are later post-processed for the performance analysis detailed in the following section.

6.7 Experimental Evaluation

The experimental evaluation demonstrates the feasibility of implementing our optimized communication scheme for infrastructure-less sensing using batteryless sensor nodes. The scheme is realizable in an application with a small energy footprint (Section 6.7.2) and scalable for dense deployment of a high number of independent sensor beacons and receivers (Section 6.7.3). By comparison of the simulation to the experimental results in Section 6.7.4 it is confirmed that the model abstracts the characteristics of the considered scenario with high accuracy. The optimized communication scheme leads to a significant improvement in terms of data utility over a simple baseline strategy as is shown in Section 6.7.5. Finally, the real-world deployment presented in Section 6.7.6 demonstrates the long-term autonomous operation of sensor beacons and the harvesting aware adaptation of the communication policy.

6.7.1 Evaluation Setups

We begin with the introduction of the setups used for all experiments and simulations. Each experiment or simulation has two different aspects: Firstly, a problem specification




Figure 6.8: Problem specification with utility specification \tilde{Q}^{spec} and a stochastic listening interval distribution \tilde{R} .

defining the required reception probability \tilde{Q}^{spec} and the receiver model \tilde{R} . Secondly, a setup which specifies how a given evaluation is conducted.

In the evaluation we use the problem specification $(\tilde{Q}^{\text{spec}}, \tilde{R})$ as plotted in Figure 6.8, and a data packet size of $N = 4$ that matches the sensor node implementation. It should however be noted that the proposed scheme is very flexible and supports arbitrary specifications of \tilde{Q}^{spec} , \tilde{R} , and N . In the considered problem specification, \tilde{Q}^{spec} specifies a decreasing importance of the 200 considered data values as a function of their age. The additional 10 aggregates represent the 10 most significant wavelet coefficients which are used for reconstruction of long-term temperature/humidity time series. The receiver active times \tilde{R} follow a Poisson distribution with arrival rate $\lambda = 30$ s. Based on this problem specification the optimized communication policies are generated for 6 power levels that correspond to one sensor node activation every $\{1, 2, 3, 4, 5, 6\}$ s.

In addition to the optimal policy, we consider a simple deterministic baseline scenario. In this scenario the most relevant data units according to \hat{Q}^{spec} are transmitted deterministically, without considering aggregate values.

The above problem specifications are evaluated in the following three setups:

Controlled Power Level Emulation In the first set of experiments the sensor node is supplied continuously from a power supply, bypassing the integrated energy management unit (EMU). The node is triggered periodically from a signal generator emulating an constant harvesting power. Upon triggering, the sensor node performs all actions part of a single activation (sense values, process and store values, transmit values, see the detailed procedure in Section 6.6.1). The activation period is set in accordance with the stated power levels. The experiment lasts for 4 h at each power level.

Energy Harvesting-Driven Execution In this experiment the sensor node is deployed with the EMU and a solar panel to operate fully autonomously. The real-world experiment is carried out in an office environment with the sensor node deployed on a table, 0.5 m away from a window. The solar panel of the transient sensor node faces towards the ceiling and is exposed to a mixture of artificial office lighting and indirect natural light. The RocketLogger measurement device introduced in Chapter 2 instruments the sensor node to trace the sensor node activations by observing the EMU trigger signal and to monitor the ambient illuminance level next to the solar panel. Illuminance levels varying from 300 lx to 800 lx at daytime and a few shorter intervals of several minutes with illuminance levels up to 1400 lx were observed during the 45 h long experiment.

Model Based Scenario Simulation The performance in terms of data utility $U(Q, Q^{\text{spec}})$ is evaluated in different scenarios using the model based simulation. It consists of the comparison of the optimized communication policy to the baseline scenario, and a model sensitivity analysis with respect to the receiver characteristics $\tilde{\mathbf{R}}$.

In the two experimental setups, an always-on smartphone is placed near the sensors and continuously logs the received data packets. The intermittent receiver behavior characterized by $\tilde{\mathbf{R}}$ is introduced in the post-processing, by only considering the packets received within the time window where the specific intermittent receivers are active. This analysis was performed independently for each power level, and we randomly generate 4000 intermittent receivers for each level. The time at which they start listening is chosen from a uniform random distribution and their listening interval is sampled from $\tilde{\mathbf{R}}$.

In the *energy harvesting-driven execution* setup, the power level is dictated by the environment. To perform the analysis in this setup, we identify time windows in which the average power remains relatively constant. Specifically, when performing the analysis for a power level with an activation interval δ , we consider 60 s periods where the average activation interval of the sensor node δ_{avg} is within $\delta \pm 0.1$. The intermittent receivers are activated such that their listening intervals are contained in the aforementioned time windows of relatively constant power.

6.7.2 Sensor Node Characterization

To calibrate the energy model, the energy required for one task activation, $E^{\text{activation}}$, is needed. The implemented sensor node was characterized using the *controlled power level emu-*

Table 6.1: Transient Bluetooth Low Energy (BLE) sensor node operation characteristics in terms of energy and time: mean of 12 819 activations, including the 95 % percentile of the absolute deviation from the mean.

	Energy/Power	Duration
<i>Activation Total</i> (3.3 V)	$162.90 \pm 1.01 \mu\text{J}$	$12.62 \pm 0.14 \text{ ms}$
<i>Deep Sleep</i> (2.3 V)	$3.91 \mu\text{W}$	-
<i>Run-Time Mechanism</i> (3.3 V)	$8.08 \pm 0.12 \mu\text{J}$	$0.73 \pm 0.01 \text{ ms}$

lation setup. The resulting energy model is then used in the offline policy pool generation, as well as at run time for input power estimation. The RocketLogger measurement device introduced in Chapter 2 was used to characterize the energy consumption of a single activation (at 3.3 V), as well as the system’s deep sleep power (at 2.3 V). To collect a significant number of samples the external triggering was kept constant at short periods. The results for more than 12 000 activations are summarized in Table 6.1 and demonstrate a very small energy requirement of only $E^{\text{activation}} \approx 163 \mu\text{J}$ per activation. Only a small amount thereof, i.e., $8 \mu\text{J}$, is needed for the communication policy run-time mechanism, demonstrating the low overhead of the presented communication scheme.

6.7.3 Model-Based Scalability Analysis

For the scalability analysis of many transmitting sensor beacons we consider M devices within communication range of each other. The wireless channel accesses are assumed random, as each sensor beacon independently harvests energy from the non-deterministic and spatially variable environment. With a time on air of 0.376 ms for a full BLE advertisement packet [SR18] and a sensor activation period of 10 s, the

Table 6.2: The model error analysis for the model accuracy experiment shows very small errors for all power levels pwr^{in} .

δ	Avg. Power [μW]	α	Max. Error	RMSE
1	163	0.77	0.0481	0.0132
2	81.5	0.58	0.0420	0.0113
3	54.33	0.53	0.0419	0.0132
4	40.75	0.5	0.0388	0.0157
5	32.6	0.47	0.0341	0.0121

collision probability for M sensor beacons operating in parallel is calculated as:

$$\Pr_{\text{coll}} = 1 - (1 - q)^{(M-1)} = 1 - \left(1 - 2 \cdot \frac{0.376 \text{ ms}}{10 \text{ s}}\right)^{(M-1)} \quad (6.14)$$

The beacons broadcast a data packet on all BLE advertisement channels to increase the chance of reception for scanning receiver devices. The collision probability \Pr_{coll} is still calculated using (6.14) as the deterministic advertisement channel hopping sequence guarantees no collision on any channel if there is no collision on the first advertisement channel. For a dense deployment of $M = 100$ beacons within the restricted communication range, equation (6.14) evaluates to a collision probability of less than 0.8%. Collisions with other sensor beacons can therefore be considered minimal when compared to the higher probability of interference from other sources like WiFi or Bluetooth communication [HSD15].

The fully passive receiver behavior guarantees by design a scalability an unbounded number of receiver devices.

6.7.4 Model Accuracy

To validate the model accuracy, the *controlled power level emulation* setup described in Section 6.7.1 was used. The




Figure 6.9: The model accuracy experiment results with $\delta = 1$ demonstrate a high model accuracy.

reception probabilities Q for the power level with $\delta = 1$ are shown in Figure 6.9. This figure compares the ideal (Q^{lossless}), the measured ($Q^{\text{experiment}}$), and the packet loss compensated reception probability (Q^{lossy}) with the scaled specification $\alpha \cdot Q^{\text{spec}}$. The scaling factor α relates to the maximum utility that resulted from the offline schedule optimization as specified in (6.3). The first Q refers to the theoretical value obtained by solving the optimization formulation introduced in Section 6.4. Comparing it to $\alpha \cdot Q^{\text{spec}}$ shows that the optimization formulation in (6.7) to (6.9) finds a correct packet selection policy ρ that results in a data reception probability Q that maximizes the data utility as specified in (6.3). The experimental results for Q are slightly lower than the model, because the model assumes a lossless communication channel. In the actual deployment the sensor nodes were placed in an office environment, where the frequency channel is heavily occupied. Augmenting the model to account for a time-invariant packet reception rate




Figure 6.10: The data utility provided by policies optimized for specific values of λ and the baseline scheme in comparison to the optimally achievable data utility.

PRR as discussed in Section 6.5.4 and using the empirically observed value (91.8% for the shown power level), results in \mathbf{Q} lossy. With this compensation factor the theoretical and the experimentally observed reception probabilities match. The model error analysis for all power levels is summarized in Table 6.2. The table states the maximum absolute error and root mean square error (RMSE) between the measured and packet loss compensated \mathbf{Q} for all power levels. The consistently low errors across all power levels demonstrate the correctness of our model introduced in Section 6.3.2, and the correct implementation of the communication scheme.

6.7.5 Data Utility Analysis

We perform a sensitivity analysis of the communication policy with respect to the receiver's listening characteristic $\bar{\mathbf{R}}$. Furthermore, we compare the proposed communication policy to the baseline policy. Both evaluations use the *model*




Figure 6.11: Comparison of the data utility of the baseline and optimal scheme for a problem specification with a history coverage shortened to 25 s.

based scenario simulation setup and do not consider aggregate values, since the baseline policy does not support sending any aggregate values.

To perform the sensitivity analysis, we examine the impact on data utility when the actual receiver behavior differs from the model used during optimization. We generate the optimal communication policies for 25 additional receiver models and calculate their data utility. In these models the receiver active times follow a Poisson distribution with arrival rates $\lambda \in \{0, 1, \dots, 25\}$.

The results of the sensitivity analysis are given in Figure 6.10. The *optimal policy* shows the utility when the actual receiver behavior (the specific value of λ) is identical to the value for which the optimal communication policy is generated. This represents the maximum achievable utility, which requires precise characterization of the receiver behavior. The *policy for $\lambda = 15$* , henceforth called *policy 15* shows the data utility for different λ of the simulated

receiver when the optimal communication policy is generated for $\lambda = 15$. The data utility of policy 15 is significantly lower than the optimum for $\lambda < 15$. This shows that the data utility is fairly sensitive to the receiver characteristics. Therefore, it is important to optimize the policy for the specific receiver behavior. For $\lambda = 15$, policy 15 provides the optimal utility, since there is no mismatch between the actual receiver behavior and the one used for generating the optimal communication policy. For $\lambda > 15$, we see that the data utility remains unchanged. Even though the receiver is listening for longer, additional gains in utility cannot be observed. This is explained by the fact that the value of the data utility is dictated by the most recent sensor values (see (6.2)). The reception probability of these values do not improve if the receiver is active for longer. Figure 6.10 also shows the data utility for communication policies optimized with $\lambda = 10, 20$. The data utility of these policies follow a similar trend as policy 15, while reaching the optimal data utility at their respective optimization points.

For comparison, we include the data utility analysis for the baseline policy. For the given \tilde{Q}^{spec} the most relevant values correspond to the most recently sampled ones. Those are the actual values selected for transmission by the baseline policy. The utility of the baseline policy is 0 for all values of λ , because it has zero probability of transmitting long-term historical values. However, the baseline policy is expected to provide an increased data utility when the receiver listening interval is similar to the relevant history window. To evaluate this aspect, we modify the problem specifications shown in Figure 6.8 by setting the \tilde{Q}^{spec} for sensor readings with an age greater than 25 s to 0, while keeping the original receiver specification. The resulting data utility for both approaches are shown in Figure 6.11, again as a function of the λ of the




Figure 6.12: Real-world harvesting use-case results with $\delta = 1$: high model accuracy despite additional sources of error.

simulated receiver. For low values of λ the baseline data utility remains at 0. This is because the receiver is active for a short amount of time and the probability of receiving sensor values with an age of 25 s is 0. However, as λ increases the data utility starts increasing and approaches the optimal data utility near $\lambda = 25$. At this point a receiver gets all important parts of the history with a probability close to 1. It is important to note that the utility achieved by the optimal communication policy is no less than the utility of the baseline policy for any value of λ . This substantiates the optimality of the proposed communication scheme.

6.7.6 Real-World Harvesting Use-Case

To evaluate the model accuracy under real-world harvesting conditions we used the *energy harvesting-driven execution* setup detailed in Section 6.7.1. The resulting reception probabilities Q are presented in Figure 6.12 for the power level

Table 6.3: The model errors analysis for real-world harvesting use-case experiment in an office deployment with illuminance levels varying from 300 lx to 800 lx demonstrates high model accuracy. The statistically not significant results for the $\delta = 2$ were omitted.

δ	Avg. Power [μ W]	α	Max. Error	RMSE
1	163	0.77	0.1099	0.0306
3	54.33	0.53	0.0875	0.0334
4	40.75	0.5	0.0970	0.0372
5	32.6	0.47	0.1099	0.0481
6	27.17	0.48	0.1457	0.0610

with $\delta = 1$. The error observed here is slightly higher when compared to the model accuracy experiment in Section 6.7.4. This is expected, because the activation interval of the sensor node is not controlled, but variable and dictated by the environment. Small deviations in sensor node activation intervals are the primary cause of the higher errors in the data reception probabilities. However, the results still show a high model accuracy.

Table 6.3 summarizes the model error analysis for the packet loss compensated model at different power levels. The results for $\delta = 2$ were omitted, because the sensor node did not operate at this power level for a statistically significant amount of time. The table shows low errors for all power levels with a maximum RMSE of 0.061 for the $\delta = 6$ power level. The high model accuracy for all power levels demonstrates that the sensor node is able to dynamically adapt the communication policy to changing harvesting conditions.

6.8 Summary

We have motivated the use of environment-powered infrastructure-less sensor beacons for autonomous long-term monitoring applications. In the considered scenario, receivers are mobile and therefore intermittently within the communication range of sensor beacons. To address the challenge of communicating useful information in such a setting, we quantified the usefulness of the sensed information using the notion of data utility. Combining this metric with a stochastic model for the receiver mobility, we formulated an optimization problem and derived an optimized communication scheme. Leveraging the energy management and design principles from Chapter 5, we developed a batteryless Bluetooth Low Energy (BLE) sensor node implementing our novel communication principle. We performed extensive experimental evaluation using the solar powered sensor node and commodity smartphones as receivers. The results demonstrated that the presented scheme is realizable with low overhead on resource constrained embedded platforms. Moreover, the real-world deployment matched the theoretically expected values with a root mean square error (RMSE) in data reception probability that was limited to 0.061 in any evaluated setup. Finally, a multi-day harvesting experiment under indoor office light conditions emphasized that the proposed scheme gracefully adapts to varying harvesting conditions.

Appendices

6.A Optimization Formulation Convexity

Here, we prove that the optimization formulation proposed in (6.7) to (6.9) and its extension with aggregate values in Section 6.5.2 is convex in \mathbf{P} . We will refer to Q_i as $Q_i(\mathbf{P}, \mathbf{R})$ to explicitly show that it is a function of \mathbf{P} and \mathbf{R} .

6.A.1 Convexity of Optimization Formulation

The optimization formulation specified in equations (6.7) to (6.9) is convex.

Lemma 6.1. *The function $d(i, \mathbf{P}) = -Q_i(\mathbf{P}, \mathbf{R})/Q_i^{spec}$ is convex $\forall i \in \{0, 1 \dots Q^{spec*}\}$.*

Proof. To prove convexity, we need to prove that $d(i, \mathbf{P})$ is twice differentiable and that its Hessian is positive-semi-definite. Differentiating $d(i, \mathbf{P})$ w.r.t P_u and P_v , we get the

following expression:

$$d_{P_u P_v}(i, \mathbf{P}) = \frac{1}{Q_i^{\text{spec}}} \cdot \begin{cases} 0 & \text{if } v = u \\ & \forall u > i \\ & \forall v > i \\ \sum_{j=i-w}^{j \leq P^*} R_j \cdot \prod_{k \in S(i,j)/\{u,v\}} (1 - P_k) & \text{otherwise} \end{cases} \quad (6.15)$$

where $w = \min\{u, v\}$. Since \mathbf{R} has no negative element and $0 \leq P_i \leq 1 \quad \forall i \in \{0, \dots, P^*\}$, $d_{P_u P_v}(\mathbf{P})$ is non-negative $\forall i, u, v \in \{0, \dots, P^*\}$. Therefore, the Hessian of $d(i, \mathbf{P})$ is positive-semi-definite in the specified domain of \mathbf{P} . Having a semi-positive-definite Hessian is a sufficient condition for convexity [BV04] and $d(i, \mathbf{P})$ is convex in \mathbf{P} . \square

Lemma 6.2. *The objective function (6.7) is convex with respect to \mathbf{P} .*

Proof. From Lemma 6.1, it can be seen that $-Q_i/Q_i^{\text{spec}}$ is convex with respect to \mathbf{P} . The objective in (6.7) can be transformed to the following dual form:

$$\text{minimize} \quad \max_{0 \leq i \leq Q^{\text{spec}*}} \{-Q_i/Q_i^{\text{spec}}\} \quad (6.16)$$

In (6.16), the max of n convex functions is computed. As the max of convex functions is also convex, the objective requires minimizing a convex function. \square

Theorem 6.1. *The optimization formulation specified in equations (6.7) to (6.9) is convex.*

Proof. The objective (6.7) requires maximizing a concave function or minimizing a convex function (Lemma 6.2). Constraints (6.8) and (6.9) are affine. Therefore, the optimization problem is convex. \square

6.A.2 Convexity of Aggregates Extension

Augmenting the optimization formulation in (6.7) to (6.9), retains its convexity.

Lemma 6.3. *The function $d^{\text{agg}}(i, \mathbf{P}) = -Q_i^{\text{agg}}/Q_i^{\text{spec,agg}}$ is convex $\forall i \in \{0, 1, \dots, Q^{\text{spec,agg}^*}\}$.*

Proof. To prove convexity, we need to prove that $d^{\text{agg}}(i, \mathbf{P})$ is twice differentiable and that its Hessian is positive-semi-definite.

$$d^{\text{agg}}_{P_u P_v}(i, \mathbf{P}) = \frac{1}{Q_i^{\text{spec}}} \cdot \begin{cases} \sum_{j=1}^{j \leq R^*} j(j+1)R_j(1-P_k)^{(j-1)} & \text{if } u = v = i \\ 0 & \text{otherwise} \end{cases} \quad (6.17)$$

Using arguments similar to the ones used in Lemma 6.1, it can be seen that the Hessian of $d^{\text{agg}}(i, \mathbf{P})$ is positive-semi-definite, proving its convexity. \square

Theorem 6.2. *The optimization formulation specified in equations (6.7) to (6.9) remains convex after adding aggregate values to the formulation.*

Proof. We have proven that the individual components of the objective for aging and aggregate values are convex (Lemmas 6.1 and 6.3). The objective (6.7) requires maximizing a concave function or minimizing a convex function

(Lemma 6.2). Constraints equations (6.8) and (6.9) are affine. Therefore, the optimization problem is convex. \square

7

Conclusions and Outlook

This chapter summarizes the contributions of this thesis and outlines possible future research directions in the area of environment-powered embedded systems.

7.1 Contributions

In this thesis, we have focused on the use of small-scale in-situ energy harvesting as an energy supply for cyber-physical systems (CPSs). Considering scenarios where the energy input is highly constrained, we aimed at exploring the potential and limitations of energy harvesting when limiting the energy buffering to tiny storage elements. In parallel, we introduced the infrastructure necessary for thorough in-situ observation of energy harvesting systems, and the emulation of the physical environment and application's electrical properties for exhaustive evaluation in the lab.

The main contributions of this thesis are presented in the

following, starting with the tools that were used to experimentally validate the subsequently presented methodologies.

In-Situ Mixed-Signal Measurements

We developed the RocketLogger, an all-in-one measurement solution offering joint observation of the environment's physical signals, of multiple energy flows, and of digital states. Integrating this broad range of measurement capabilities into a portable device significantly enhances the in-situ observability of the multi-faceted aspects of energy harvesting systems. The RocketLogger proved to be a highly valuable tool for thorough in-situ validation of the methodologies presented throughout this thesis.

Environment and Application Testbed

We introduced a testbed to consistently reproduce the physical properties of photovoltaic and thermoelectric energy harvesting, as well as a SmartLoad for time- and event-triggered emulation of energy consumption profiles of typical low-power systems. Centralized control over these main sources of non-determinism in energy harvesting systems enables rapid and thorough evaluation of energy harvesting systems with an unprecedented coverage of operating conditions. Controlled speed-up of setpoint characterizations and the accelerated replay of real-world traces provide timely feedback in the design process of energy harvesting systems. The testbed was shown to be highly valuable for the implementation and evaluation of the methodologies summarized in the following.

Thermoelectric Harvesting the Ground-to-Air Interface

We have presented an efficient thermoelectric harvesting platform for extracting electrical energy from small bipolar thermal gradients occurring at the ground-to-air boundary. Introducing the first model that incorporates all aspects from the environment to the application, we significantly extended the modeling of the physical aspects of energy harvesting cyber-physical systems (CPSs). In combination with a newly proposed rectifier circuit, an optimized energy harvesting system was implemented and extensively evaluated in real-world deployments. The evaluation attests the accuracy of the model and demonstrates an unprecedented output power in the investigated harvesting scenario. The platform harvested up to 27.2 mW in direct sunlight and 6.3 mW during nighttime and considerably outperformed the state-of-the-art both in terms of average and maximum power.

Efficient Energy Management for Batteryless Systems

We proposed a novel energy management principle that decouples the efficient extraction of energy from the non-deterministic environment and the energy supply requirements of the load. By enabling both ends to operate at their respective optimal operating point, the application performance can be maximized. This principle is implemented in an energy management unit (EMU) that performs maximum power point tracking (MPPT) on the input, and exposes an interface to enable a load to dynamically configure its operating point. This guarantees high efficiency and application performance. At the same time the presented method mitigates the variability of the energy input using a minimized energy buffer. Extensive evaluation of this principle demonstrated the feasibility of supplying applications reliably and with high efficiency. This holds even under very low or highly variable

energy input and power requirements of the application that are 3 orders of magnitude higher than the input.

Infrastructure-Less Monitoring with Batteryless Sensors

Lastly, we motivated the use of batteryless sensor beacons for long-term infrastructure-less monitoring applications. We studied the utility of data transmissions in such a communication scenario. Based on a model of the scenario, we derived a probabilistic data transmission scheme that optimizes the data utility of opportunistically listening receivers. We leveraged the energy management principle introduced previously to implement this communication scheme on batteryless Bluetooth Low Energy (BLE) sensor nodes. The results from extensive real-world experiments closely matched the theoretical simulations and demonstrated that the scheme gracefully adapts to variable harvesting conditions.

7.2 Future Directions

In the presented thesis we developed novel methods that enable the design of energy harvesting systems to operate reliably under very low and highly variable energy input. More so, we demonstrated that these methods are implemented with low energy overhead and therefore enable systems to operate efficiently under these conditions. We claim that these methods provide an important building block for the design of reliable batteryless applications.

Application Level Performance Guarantees

While the methods presented provide reliable execution of the application, the rate at which it executes is still dictated

by the non-deterministic environment. This makes providing application level performance guarantees difficult. To increase the robustness against the environment variability, the diverse physical signals exploitable for energy harvesting can be used to mitigate the variability of each other. This could be performed at a node level by harvesting from multiple sources. Alternatively, multiple nodes that sense the same quantities, but employ different harvesting modalities, could offer a more robust performance as a group. In such a distributed sensing approach, coordination across nodes is required in order to collectively provide a high quality of service.

Batteryless Networking

To enable communication among sensor nodes for means of coordinated and/or distributed sensing demands for bi-directional communication. As batteryless sensor nodes are strongly constrained in terms of energy, tight time synchronization among nodes is required to precisely schedule the infrequent and short radio activations for reliable communication. Efficient bootstrapping and maintaining such a time synchronization present major challenges, in particular under the highly sporadic operation of batteryless systems. It is possible that specific signatures of the environment coinciding with energy harvesting, such as sunrise or switching on a light, can be leveraged for synchronization purposes. Continuously maintaining synchronization requires highly accurate time-keeping and frequent resynchronization with other nodes to not fall back to costly bootstrapping. On a sensor node level the maintenance of such time-keeping and networking services demands additional energy management mechanisms to satisfy the multiple and potentially conflicting application requirements. These aspects illustrate that there are many innovative solutions required to enable networking

with batteryless systems.

Comparison of Harvesting Modalities

Not directly related to the previous directions, but a relevant problem for assessing the performance of energy harvesting, is the comparison of different harvesting modalities exploiting the same harvesting source. As was discussed in Chapter 4, a systematic comparison of harvesting modalities such as thermoelectric generator (TEG) and photovoltaic cells is highly challenging due to the large number of factors affecting their efficiency. A basic comparison, for example by area, does not suffice, as further aspects including spectral properties, converter architectures, illumination or temperature impact the harvesting efficiency. Many of these factors also vary dependent on the exact location at which the systems are deployed or exhibit a significant time dependence. The development of a methodology that enables such a comparison demands detailed exploration of requirements and of parameters to consider.

Bibliography

- [ABB⁺17] L. V. Allmen, G. Bailleul, T. Becker, J.-D. D. Decotignie, M. E. Kiziroglou, C. Leroux, P. D. Mitcheson, J. Muller, D. Piguet, T. T. Toh, A. Weisser, S. W. Wright, and E. M. Yeatman. Aircraft Strain WSN Powered by Heat Storage Harvesting. *IEEE Transactions on Industrial Electronics*, 64(9):7284–7292, September 2017. doi:10.1109/TIE.2017.2652375.
- [ABD⁺19] R. Ahmed, B. Buchli, S. Draskovic, L. Sigrist, P. Kumar, and L. Thiele. Optimal Power Management with Guaranteed Minimum Energy Utilization for Solar Energy Harvesting Systems. *ACM Transactions on Embedded Computing Systems*, 18(4):1–26, June 2019. doi:10.1145/3317679.
- [AGCCLS16] F. Attivissimo, C. Guarnieri Calò Carducci, A. Lanzolla, and M. Spadavecchia. An Extensive Unified Thermo-Electric Module Characterization Method. *Sensors*, 16(12):2114, December 2016. doi:10.3390/s16122114.
- [AH09] J. Andersen and M. T. Hansen. Energy Bucket: A Tool for Power Profiling and Debugging of Sensor Nodes. In *Proceedings of the 3rd International Conference on Sensor Technologies and Applications (SENSORCOMM)*, pages 132–138, Athens, Greece, June 2009. IEEE. doi:10.1109/SENSORCOMM.2009.29.

- [AM03] L. I. Anatyckuk and P. D. Mikityuk. Thermal Generators Using Heat Flows in Soils. In *Proceedings of the 22nd International Conference on Thermoelectrics (ICT)*, pages 598–601. IEEE, 2003. doi:10.1109/ICT.2003.1287584.
- [AM16] K. Ahmed and S. Mukhopadhyay. A 190 nA Bias Current 10 mV Input Multistage Boost Regulator With Intermediate-Node Control to Supply RF Blocks in Self-Powered Wireless Sensors. *IEEE Transactions on Power Electronics*, 31(2):1322–1333, February 2016. doi:10.1109/TPEL.2015.2423666.
- [BASM16] N. A. Bhatti, M. H. Alizai, A. A. Syed, and L. Motola. Energy Harvesting and Wireless Transfer in Sensor Network Applications: Concepts and Experiences. *ACM Transactions on Sensor Networks*, 12(3):24:1–24:40, August 2016. doi:10.1145/2915918.
- [BBF⁺11] J. Beutel, B. Buchli, F. Ferrari, M. Keller, M. Zimmerling, and L. Thiele. X-SENSE: Sensing in Extreme Environments. In *Proceedings of the 2011 Design, Automation and Test in Europe Conference and Exhibition (DATE)*, pages 1–6, March 2011. doi:10.1109/DATE.2011.5763236.
- [BBMT08] D. Brunelli, L. Benini, C. Moser, and L. Thiele. An Efficient Solar Energy Harvester for Wireless Sensor Nodes. In *Proceedings of the 2011 Design, Automation and Test in Europe Conference and Exhibition (DATE)*, pages 104–109, 2008. doi:10.1109/DATE.2008.4484670.
- [BBR⁺15] S. Bobovych, N. Banerjee, R. Robucci, J. P. Parkerson, J. Schmandt, and C. Patel. SunaPlayer: High-accuracy Emulation of Solar Cells. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 59–70. ACM, 2015. doi:10.1145/2737095.2737110.
- [BC12] S. Bandyopadhyay and A. P. Chandrakasan. Platform Architecture for Solar, Thermal, and Vibration

- Energy Combining With MPPT and Single Inductor. *IEEE Journal of Solid-State Circuits*, 47(9):2199–2215, September 2012. doi:10.1109/JSSC.2012.2197239.
- [BdM00] L. Benini and G. de Micheli. System-Level Power Optimization: Techniques and Tools. *ACM Transactions on Design Automation of Electronic Systems*, 5(2):115–192, April 2000. doi:10.1145/335043.335044.
- [Bea] BeagleBoard.org Foundation. SeeedStudio Beagle-Bone Green. <https://beagleboard.org/green/>. last accessed 12.01.2020.
- [BGH⁺09] J. Beutel, S. Gruber, A. Hasler, R. Lim, A. Meier, C. Plessl, I. Talzi, L. Thiele, C. Tschudin, M. Woehrle, others, and M. Yuecel. PermaDAQ: A Scientific Instrument for Precision Sensing and Data Recovery in Environmental Extremes. In *Proceedings of the 8th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 265–276. IEEE, April 2009.
- [BHE00] N. Bulusu, J. Heidemann, and D. Estrin. GPS-Less Low-Cost Outdoor Localization for Very Small Devices. *IEEE Personal Communications*, 7(5):28–34, Oct./2000. doi:10.1109/98.878533.
- [BKS⁺12] A. M. Bazzi, Z. Klein, M. Sweeney, K. P. Kroeger, P. S. Shenoy, and P. T. Krein. Solid-State Solar Simulator. *IEEE Transactions on Industry Applications*, 48(4):1195–1202, July 2012. doi:10.1109/TIA.2012.2199071.
- [BM17] N. A. Bhatti and L. Mottola. HarvOS: Efficient Code Instrumentation for Transiently-Powered Embedded Sensing. In *Proceedings of the 16th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 209–219, 2017. doi:10.1145/3055031.3055082.
- [BMAS19] A. Branco, L. Mottola, M. H. Alizai, and J. H. Siddiqui. Intermittent Asynchronous Peripheral Operations. In

- Proceedings of the 17th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, page 13, 2019. doi:10.1145/3356250.3360033.
- [BMG18] D. Blalock, S. Madden, and J. Guttag. Sprintz: Time Series Compression for the Internet of Things. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–23, September 2018. doi:10.1145/3264903.
- [Bru75] W. Brutsaert. On a Derivable Formula for Long-Wave Radiation from Clear Skies. *Water Resources Research*, 11(5):742–744, 1975. doi:10.1029/WR011i005p00742.
- [BSBT14a] B. Buchli, F. Sutton, J. Beutel, and L. 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)*, pages 31–45, New York, NY, USA, 2014. ACM. doi:10.1145/2668332.2668333.
- [BSBT14b] B. Buchli, F. Sutton, J. Beutel, and L. Thiele. Towards Enabling Uninterrupted Long-Term Operation of Solar Energy Harvesting Embedded Systems. In *Proceedings of the 11th European Conference on Wireless Sensor Networks (EWSN 2014)*, pages 66–83. Springer, February 2014. doi:10.1007/978-3-319-04651-8_5.
- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004. doi:10.1017/CBO9780511804441.
- [BWD⁺16] D. Balsamo, A. S. Weddell, A. Das, A. R. Arreola, D. Brunelli, B. M. Al-Hashimi, G. V. Merrett, and L. Benini. Hibernus++: A Self-Calibrating and Adaptive System for Transiently-Powered Embedded Devices. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 35(12):1968–1980, 2016. doi:10.1109/TCAD.2016.2547919.

-
- [BWM⁺15] D. Balsamo, A. S. Weddell, G. V. Merrett, B. M. Alhashimi, D. Brunelli, and L. Benini. Hibernus: Sustaining Computation during Intermittent Supply for Energy-Harvesting Systems. *IEEE Embedded Systems Letters*, 7(1):1–4, 2015. doi:10.1109/LES.2014.2371494.
- [BZL14] N. Brouwers, M. Zuniga, and K. Langendoen. NEAT: A Novel Energy Analysis Toolkit for Free-Roaming Smartphones. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems (SenSys)*, pages 16–30, New York, NY, USA, 2014. ACM. doi:10.1145/2668332.2668337.
- [CK18] Y. Collet and M. Kucherawy. Zstandard Compression and the application/zstd Media Type. Technical Report RFC8478, RFC Editor, October 2018.
- [CKL00] N. Chang, K. Kim, and H. G. Lee. Cycle-Accurate Energy Consumption Measurement and Analysis: Case Study of ARM7TDMI. In *Proceedings of the 2000 International Symposium on Low Power Electronics and Design (ISLPED)*, 2000. doi:10.1109/LPE.2000.155275.
- [CL16] A. Colin and B. Lucia. Chain: Tasks and Channels for Reliable Intermittent Programs. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 514–530. ACM, 2016. doi:10.1145/2983990.2983995.
- [CRL18] A. Colin, E. Ruppel, and B. Lucia. A Reconfigurable Energy Storage Architecture for Energy-Harvesting Devices. In *Proceedings of the 23rd International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, ASPLOS '18, pages 767–781. ACM, 2018. doi:10.1145/3173162.3173210.
- [CW08] E. Candes and M. Wakin. An Introduction To Compressive Sampling. *IEEE Signal Processing Magazine*, 25(2):21–30, March 2008. doi:10.1109/MSP.2007.914731.

- [DAS08] C. Domínguez, I. Antón, and G. Sala. Solar Simulator for Concentrator Photovoltaic Systems. *Optics Express*, 16(19):14894, September 2008. doi:10.1364/OE.16.014894.
- [DBL⁺15] E. Dallago, A. L. Barnabei, A. Liberale, P. Malcovati, and G. Venchi. An Interface Circuit for Low-Voltage Low-Current Energy Harvesting Systems. *IEEE Transactions on Power Electronics*, 30(3):1411–1420, March 2015. doi:10.1109/TPEL.2014.2322521.
- [DDP18] U. Datta, S. Dessouky, and A. T. Papagiannakis. Thermal Energy Harvesting from Asphalt Roadway Pavement. In L. Mohammad, editor, *Advancement in the Design and Performance of Sustainable Asphalt Pavements*, pages 272–286. Springer International Publishing, 2018. doi:10.1007/978-3-319-61908-8_20.
- [DFG⁺09] S. Dalola, V. Ferrari, M. Guizzetti, D. Marioli, E. Sardini, M. Serpelloni, and A. Taroni. Autonomous Sensor System with Power Harvesting for Telemetric Temperature Measurements of Pipes. *IEEE Transactions on Instrumentation and Measurement*, 58(5):1471–1478, May 2009. doi:10.1109/TIM.2009.2012946.
- [DFPC08] P. Dutta, M. Feldmeier, J. Paradiso, and D. Culler. Energy Metering for Free: Augmenting Switching Regulators for Real-Time Monitoring. In *Proceedings of the 7th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 283–294. IEEE, April 2008. doi:10.1109/IPSN.2008.58.
- [DGV04] A. Dunkels, B. Gronvall, and T. Voigt. Contiki - a Lightweight and Flexible Operating System for Tiny Networked Sensors. In *Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks*, pages 455–462, November 2004. doi:10.1109/LCN.2004.38.
- [DLBL⁺16] E. Dallago, A. Lazzarini Barnabei, A. Liberale, G. Torelli, and G. Venchi. A 300-mV Low-Power

- Management System for Energy Harvesting Applications. *IEEE Transactions on Power Electronics*, 31(3):2273–2281, March 2016. doi:10.1109/TPEL.2015.2431439.
- [DQC⁺17] F. Deng, H. Qiu, J. Chen, L. Wang, and B. Wang. Wearable Thermoelectric Power Generators Combined with Flexible Supercapacitor for Low-Power Human Diagnosis Devices. *IEEE Transactions on Industrial Electronics*, 64(2):1477–1485, February 2017. doi:10.1109/TIE.2016.2613063.
- [FdJ19] R. P. D. A. Filho, O. A. de Lima Junior, and C. G. F. Junior. An FPGA-Based Evaluation Platform for Energy Harvesting Embedded Systems. In *Proceedings of the 32nd Symposium on Integrated Circuits and Systems Design (SBCCI)*, pages 1–6, São Paulo, Brazil, 2019. ACM. doi:10.1145/3338852.3339863.
- [FPSV05] N. Femia, G. Petrone, G. Spagnuolo, and M. Vitelli. Optimization of Perturb and Observe Maximum Power Point Tracking Method. *IEEE Transactions on Power Electronics*, 20(4):963–973, July 2005. doi:10.1109/TPEL.2005.850975.
- [FY18] S. Feng and J. Yang. Optimal Status Updating for an Energy Harvesting Sensor with a Noisy Channel. In *Proceedings of the 2018 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 348–353, Honolulu, HI, April 2018. IEEE. doi:10.1109/INFOCOMW.2018.8406974.
- [GCZ19] K. Geissdoerfer, M. Chwalisz, and M. Zimmerling. Shepherd: A Portable Testbed for the Batteryless IoT. In *Proceedings of the 17th ACM Conference on Embedded Network Sensor Systems (SenSys)*, pages 83–95. ACM, 2019. doi:10.1145/3356250.3360042.
- [GDHE⁺20] M. A. Green, E. D. Dunlop, J. Hohl-Ebinger, M. Yoshita, N. Kopidakis, and A. W. Ho-Baillie. Solar Cell Efficiency Tables (Version 55). *Progress in*

- Photovoltaics: Research and Applications*, 28(1):3–15, January 2020. doi:10.1002/pip.3228.
- [GM74] W. Gray and R. Müller. Direct Radiative Transfer. In *Engineering Calculations in Radiative Heat Transfer*, pages 26–45. Elsevier, 1974. doi:10.1016/B978-0-08-017787-8.50007-8.
- [GPB⁺15] A. Gomez, C. Pinto, A. Bartolini, D. Rossi, L. Benini, H. Fatemi, and J. P. de Gyvez. Reducing Energy Consumption in Microcontroller-Based Platforms with Low Design Margin Co-Processors. In *Proceedings of the 2015 Design, Automation Test in Europe Conference and Exhibition (DATE)*, pages 269–272. IEEE, 2015. doi:10.7873/DATE.2015.1013.
- [GPH⁺03] S. Gruber, M. Peter, M. Hoelzle, I. Woodhatch, and W. Haerberli. Surface Temperatures in Steep Alpine Rock Faces - A Strategy for Regional-Scale Measurement and Modelling. *Proceedings of the 8th International Conference on Permafrost*, 1:325–330, 2003.
- [GSC⁺19] M. Gao, C. Su, J. Cong, F. Yang, Y. Wang, and P. Wang. Harvesting Thermoelectric Energy from Railway Track. *Energy*, 180:315–329, August 2019. doi:10.1016/j.energy.2019.05.087.
- [GSS⁺17] A. Gomez, L. Sigrist, T. Schalch, L. Benini, and L. Thiele. Wearable, Energy-Opportunistic Vision Sensing for Walking Speed Estimation. In *Proceedings of the 2017 IEEE Sensors Applications Symposium (SAS)*, pages 1–6, Glassboro, NJ, USA, March 2017. IEEE. doi:10.1109/SAS.2017.7894074.
- [GWZ13] M. Gorlatova, A. Wallwater, and G. Zussman. Networking Low-Power Energy Harvesting Devices: Measurements and Algorithms. *IEEE Transactions on Mobile Computing*, 12(9):1853–1865, September 2013. doi:10.1109/TMC.2012.154.

-
- [HD16] B. H. Hamadani and B. Dougherty. Solar Cell Characterization. In M. P. Paranthaman, W. Wong-Ng, and R. N. Bhattacharya, editors, *Semiconductor Materials for Solar Photovoltaic Cells*, volume 218, pages 229–245. Springer International Publishing, Cham, 2016. doi:10.1007/978-3-319-20331-7_8.
- [HFdB17] P. A. Hager, H. Fatemi, J. P. de Gyvez, and L. Benini. A Scan-Chain Based State Retention Methodology for IoT Processors Pperating on Intermittent Energy. In *Proceedings of the 2017 Design, Automation Test in Europe Conference and Exhibition (DATE)*, pages 1171–1176. IEEE, March 2017. doi:10.23919/DATED.2017.7927166.
- [HGH11] A. Hasler, S. Gruber, and W. Haeberli. Temperature Variability and Offset in Steep Alpine Rock and Ice Faces. *The Cryosphere*, 5(4):977–988, November 2011. doi:10.5194/tc-5-977-2011.
- [HH15] P. Horowitz and W. Hill. *The Art of Electronics*. Cambridge University Press, New York, NY, USA, 3rd edition, 2015.
- [Hic17] M. Hicks. Clank: Architectural Support for Intermittent Computation. In *Proceedings of the 44th Annual International Symposium on Computer Architecture (ISCA)*, pages 228–240, Toronto, ON, Canada, 2017. ACM. doi:10.1145/3079856.3080238.
- [HRB⁺17] L. Hanschke, C. Renner, J. Brockmann, T. Hamann, J. Peschel, A. Schell, and A. Sowarka. Light Insight - Emulation of Radiation Traces for Analysis and Evaluation of Solar-Harvesting Algorithms. In *Proceedings of the 5th International Workshop on Energy Neutral Sensing Systems (ENSsys)*. ACM, November 2017. doi:10.1145/3142992.3142994.
- [HSD15] A. Hithnawi, H. Shafagh, and S. Duquennoy. TIIM: Technology-Independent Interference Mitigation for Low-Power Wireless Networks. In *Proceedings*

- of the 14th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–12. ACM, 2015. doi:10.1145/2737095.2737104.
- [HSS12] S. S. Hashemi, M. Sawan, and Y. Savaria. A High-Efficiency Low-Voltage CMOS Rectifier for Harvesting Energy in Implantable Devices. *IEEE Transactions on Biomedical Circuits and Systems*, 6(4):326–335, August 2012. doi:10.1109/TBCAS.2011.2177267.
- [HSS15] J. Hester, L. Sitanayah, and J. Sorber. Tragedy of the Coulombs: Federating Energy Storage for Tiny, Intermittently-Powered Sensors. In *Proceedings of the 13th ACM Conference on Embedded Network Sensor Systems (SenSys)*, pages 5–16. ACM, 2015. doi:10.1145/2809695.2809707.
- [HSS17] J. Hester, K. Storer, and J. Sorber. Timely Execution on Intermittently Powered Batteryless Sensors. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems (SenSys)*, pages 1–13. ACM, 2017. doi:10.1145/3131672.3131673.
- [HSSS17] J. Hester, L. Sitanayah, T. Scott, and J. Sorber. Realistic and Repeatable Emulation of Energy Harvesting Environments. *ACM Transactions on Sensor Networks*, 13(2), April 2017. doi:10.1145/3064839.
- [HTR+16] J. Hester, N. Tobias, A. Rahmati, L. Sitanayah, D. Holcomb, K. Fu, W. P. Burleson, and J. Sorber. Persistent Clocks for Batteryless Sensing Devices. *ACM Transactions on Embedded Computing Systems*, 15(4):77:1–77:28, August 2016. doi:10.1145/2903140.
- [IAJCS17] C. Izidoro, O. Ando Junior, J. Carmo, and L. Schaeffer. Characterization of Thermoelectric Generator for Energy Harvesting. *Measurement*, 106:283–290, August 2017. doi:10.1016/j.measurement.2016.01.010.
- [JDCS07] X. Jiang, P. Dutta, D. Culler, and I. 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 (IPSN)*, pages 186–195. ACM, 2007. doi:10.1145/1236360.1236386.
- [JOW⁺02] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. S. Peh, and D. Rubenstein. Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet. *ACM SIGARCH Computer Architecture News*, 30(5):96–107, December 2002. doi:10.1145/635506.605408.
- [JRLR15] H. Jayakumar, A. Raha, W. S. Lee, and V. Raghunathan. QuickRecall: A HW/SW Approach for Computing Across Power Cycles in Transiently Powered Computers. *ACM Journal on Emerging Technologies in Computing Systems*, 12(1):1–19, August 2015. doi:10.1145/2700249.
- [JWK⁺13] W. Jiang, S. Wang, S. Kan, L. Xie, and S. Hashimoto. Digitally Controlled Multi-Phase Electronic Current Sink. In *Proceedings of the 10th IEEE International Conference on Power Electronics and Drive Systems (PEDS)*, pages 862–866. IEEE, April 2013. doi:10.1109/PEDS.2013.6527138.
- [KAS⁺16] J. E. Kim, T. Abdelzaher, L. Sha, A. Bar-Noy, and R. Hobbs. Sporadic Decision-Centric Data Scheduling with Normally-Off Sensors. In *Proceedings of the 2016 IEEE Real-Time Systems Symposium (RTSS)*, pages 135–145. IEEE, November 2016. doi:10.1109/RTSS.2016.022.
- [Kaz07] M. Kazerani. A High-Performance Controllable DC Load. In *Proceedings of the 2007 IEEE International Symposium on Industrial Electronics*, pages 1015–1020. IEEE, June 2007. doi:10.1109/ISIE.2007.4374737.
- [KCWP10] Y. Kim, N. Chang, Y. Wang, and M. Pedram. Maximum Power Transfer Tracking for a Photovoltaic-Supercapacitor Energy System. In *Proceedings of*

- the 16th ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED)*, pages 307–312, 2010. doi:10.1145/1840845.1840909.
- [Kei13a] Keithley Instruments. Keithley 2001/2002: 7 1/2- and 8 1/2-Digit High Performance Multimeter, 2013. last accessed: 12.01.2020.
- [Kei13b] Keithley Instruments. *Low Level Measurements Handbook - Precision DC Current, Voltage, and Resistance Measurements*. Keithley Instruments, 7 edition, 2013. last accessed: 12.01.2020.
- [Kei19] Keithley Instruments. 2450 SourceMeter® SMU Instrument - Datasheet, 2019. last accessed: 12.01.2020.
- [KER⁺10] T. M. Kamel, R. Elfrink, M. Renaud, D. Hohlfeld, M. Goedbloed, C. De Nooijer, M. Jambunathan, and R. Van Schaijk. Modeling and Characterization of MEMS-Based Piezoelectric Harvesting Devices. *Journal of Micromechanics and Microengineering*, 20(10):105023, October 2010. doi:10.1088/0960-1317/20/10/105023.
- [Key19] Keysight Technologies. Digital Multimeters 34460A, 34461A, 34465A (6 1/2 Digit), 34470A (7 1/2 Digit) Data Sheet, 2019. last accessed: 12.01.2020.
- [KH13] J. D. Kelly and J. D. Hedengren. A Steady-State Detection (SSD) Algorithm to Detect Non-Stationary Drifts in Processes. *Journal of Process Control*, 23(3):326–331, March 2013. doi:10.1016/j.jprocont.2012.12.001.
- [KHZS07] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava. Power Management in Energy Harvesting Sensor Networks. *ACM Transactions on Embedded Computing Systems*, 6(4):32–es, September 2007. doi:10.1145/1274858.1274870.
- [KLYK09] R.-Y. Kim, J.-S. Lai, B. York, and A. Koran. Analysis and Design of Maximum Power Point Tracking Scheme for Thermoelectric Battery Energy

- Storage System. *IEEE Transactions on Industrial Electronics*, 56(9):3709–3716, September 2009. doi:10.1109/TIE.2009.2025717.
- [KLYL17] J. Kim, S.-T. Lee, S. Yang, and J. Lee. Implementation of Thermal-Energy-Harvesting Technology on Pavement. *Journal of Testing and Evaluation*, 45(2):582–590, 2017. doi:10.1520/JTE20140396.
- [KPK12] S.-G. Kim, S. Priya, and I. Kanno. Piezoelectric MEMS for Energy Harvesting. *MRS Bulletin*, 37(11):1039–1050, November 2012. doi:10.1557/mrs.2012.275.
- [KWT⁺14] M. E. Kiziroglou, S. W. Wright, T. T. Toh, P. D. Mitcheson, T. Becker, and E. M. Yeatman. Design and Fabrication of Heat Storage Thermoelectric Harvesting Devices. *IEEE Transactions on Industrial Electronics*, 61(1):302–309, January 2014. doi:10.1109/TIE.2013.2257140.
- [LC15] H. G. Lee and N. Chang. Powering the IoT: Storage-Less and Converter-Less Energy Harvesting. In *Proceedings of the 20th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 124–129. IEEE, January 2015. doi:10.1109/ASPDAC.2015.7058992.
- [Leo11] V. Leonov. Energy Harvesting for Self-Powered Wearable Devices. In A. Bonfiglio and D. De Rossi, editors, *Wearable Monitoring Systems*, pages 27–49. Springer US, Boston, MA, 2011. doi:10.1007/978-1-4419-7384-9_2.
- [LFD⁺17] E. Lattanzi, V. Freschi, M. Dromedari, L. S. Lorello, R. Peruzzini, and A. Bogliolo. A Fast and Accurate Energy Source Emulator for Wireless Sensor Networks. *EURASIP Journal on Embedded Systems*, 2016(1):18, January 2017. doi:10.1186/s13639-016-0055-5.
- [LFZ⁺13] R. Lim, F. Ferrari, M. Zimmerling, C. Walser, P. Sommer, and J. Beutel. FlockLab: A Testbed for Distributed, Synchronized Tracing and Profiling of Wire-

- less Embedded Systems. In *Proceedings of the 12th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 153–165. ACM, April 2013. doi:10.1145/2461381.2461402.
- [LMBAG10] M. Lossec, B. Multon, H. Ben Ahmed, and C. Goupil. Thermoelectric Generator Placed on the Human Body: System Modeling and Energy Conversion Improvements. *European Physical Journal: Applied Physics*, 52(1):11103, October 2010. doi:10.1051/ep-jap/2010121.
- [Lor03] R. Lorenz. Subsurface Ambient Thermoelectric Power for Moles and Penetrators. In *Proceedings of the 2003 IEEE Aerospace Conference*, volume 2, pages 2_637–2_642. IEEE, 2003. doi:10.1109/AERO.2003.1235473.
- [LPRR10] C. Lu, S. P. Park, V. Raghunathan, and K. Roy. Efficient Power Conversion for Ultra Low Voltage Micro Scale Energy Transducers. In *Proceedings of the 2010 Design, Automation & Test in Europe Conference and Exhibition (DATE)*, pages 1602–1607. IEEE, March 2010. doi:10.1109/DATE.2010.5457066.
- [LS02] E. Lawrence and G. Snyder. A Study of Heat Sink Performance in Air and Soil for Use in a Thermoelectric Energy Harvesting Device. In *Proceedings of the 21st International Conference on Thermoelectrics (ICT)*, pages 446–449. IEEE, 2002. doi:10.1109/ICT.2002.1190357.
- [LTLC16] G. T. Le, T. V. Tran, H.-S. Lee, and W.-Y. Chung. Long-Range Batteryless RF Sensor for Monitoring the Freshness of Packaged Vegetables. *Sensors and Actuators A: Physical*, 237:20–28, January 2016. doi:10.1016/j.sna.2015.11.013.
- [LYQ11] M. Liu, Y. Yang, and Z. Qin. A Survey of Routing Protocols and Simulations in Delay-Tolerant Networks. In Y. Cheng, D. Y. Eun, Z. Qin, M. Song, and

- K. Xing, editors, *Wireless Algorithms, Systems, and Applications*, volume 6843, pages 243–253. Springer, Berlin, Heidelberg, 2011. doi:10.1007/978-3-642-23490-3_22.
- [MBS⁺16] M. Magno, D. Brunelli, L. Sigrist, R. Andri, L. Cavigelli, A. Gomez, and L. Benini. InfiniTime: Multi-Sensor Wearable Bracelet with Human Body Harvesting. *Sustainable Computing: Informatics and Systems*, 11:38–49, September 2016. doi:10.1016/j.suscom.2016.05.003.
- [MCL17] K. Maeng, A. Colin, and B. Lucia. Alpaca: Intermittent Execution Without Checkpoints. *Proceedings of the ACM on Programming Languages*, 1(OOPSLA):1–30, October 2017. doi:10.1145/3133920.
- [MCS12] P. Martin, Z. Charbiwala, and M. Srivastava. DoubleDip: Leveraging Thermoelectric Harvesting for Low Power Monitoring of Sporadic Water Use. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems (SenSys)*, pages 225–238. ACM, 2012. doi:10.1145/2426656.2426679.
- [MEK⁺12] A. Moser, M. Erd, M. Kostic, K. Cobry, M. Kroener, and P. Woias. Thermoelectric Energy Harvesting from Transient Ambient Temperature Gradients. *Journal of Electronic Materials*, 41(6):1653–1661, June 2012. doi:10.1007/s11664-011-1894-4.
- [MEK⁺16] M. Masoudinejad, J. Emmerich, D. Kossmann, A. Riesner, M. Roidl, and M. ten Hompel. A Measurement Platform for Photovoltaic Performance Analysis in Environments with Ultra-Low Energy Harvesting Potential. *Sustainable Cities and Society*, 25:74–81, August 2016. doi:10.1016/j.scs.2015.09.006.
- [Mer05] R. J. E. Merry. Wavelet Theory and Applications - A Literature Study. Technical Report DCT 2005.53, Eindhoven University of Technology, Eindhoven, 2005.

- [Mik03] P. D. Mikiyuk. Thermoelectric conversion of thermal flows in soil. *Journal of Thermoelectricity*, 1(1):84–89, 2003.
- [MKK⁺15] R. Margolies, P. Kinget, I. Kymissis, G. Zussman, M. Gorlatova, J. Sarik, G. Stanje, J. Zhu, P. Miller, M. Szczodrak, B. Vignraham, and L. Carloni. Energy-Harvesting Active Networked Tags (EnHANTs): Prototyping and Experimentation. *ACM Transactions on Sensor Networks*, 11(4):1–27, November 2015. doi:10.1145/2831236.
- [MLS⁺16] K. Ma, X. Li, K. Swaminathan, Y. Zheng, S. Li, Y. Liu, Y. Xie, J. J. Sampson, and V. Narayanan. Nonvolatile Processor Architectures: Efficient, Reliable Progress with Unstable Power. *IEEE Micro*, 36(3):72–83, May 2016. doi:10.1109/MM.2016.35.
- [MMB⁺12] M. Magno, S. Marinkovic, D. Brunelli, E. Popovici, B. O’Flynn, and L. Benini. Smart Power Unit with Ultra Low Power Radio Trigger Capabilities for Wireless Sensor Networks. In *Proceedings of the 2012 Design, Automation and Test in Europe Conference and Exhibition (DATE)*, pages 75–80. IEEE, March 2012. doi:10.1109/DATE.2012.6176436.
- [MMAA05] M. Minami, T. Morito, H. Morikawa, and T. Aoyama. Solar Biscuit: A Battery-Less Wireless Sensor Network System for Environmental Monitoring Applications. In *Proceedings of the 2nd International Workshop on Networked Sensing Systems (INSS)*, page 2007. Citeseer, June 2005.
- [MSG⁺19] M. Magno, L. Sigrist, A. Gomez, L. Cavigelli, A. Libri, E. Popovici, and L. Benini. SmarTEG: An Autonomous Wireless Sensor Node for High Accuracy Accelerometer-Based Monitoring. *Sensors*, 19(12):2747, June 2019. doi:10.3390/s19122747.
- [MSS05] Y. Meydbray, R. Singh, and A. Shakouri. Thermoelectric Module Construction for Low Temperature Gra-

- dient Power Generation. In *Proceedings of the 24th International Conference on Thermoelectrics (ICT)*, pages 348–351. IEEE, 2005. doi:10.1109/ICT.2005.1519958.
- [MTBB10] C. Moser, L. Thiele, D. Brunelli, and L. Benini. Adaptive Power Management for Environmentally Powered Systems. *IEEE Transactions on Computers*, 59(4):478–491, April 2010. doi:10.1109/TC.2009.158.
- [MYR⁺08] P. D. Mitcheson, E. M. Yeatman, G. K. Rao, A. S. Holmes, and T. C. Green. Energy Harvesting From Human and Machine Motion for Wireless Electronic Devices. *Proceedings of the IEEE*, 96(9):1457–1486, September 2008. doi:10.1109/JPROC.2008.927494.
- [NPK⁺15] S. Naderiparizi, A. N. Parks, Z. Kapetanovic, B. Ransford, and J. R. Smith. WISPCam: A Battery-Free RFID Camera. In *Proceedings of the 2015 IEEE International Conference on RFID (RFID)*, pages 166–173. IEEE, April 2015. doi:10.1109/RFID.2015.7113088.
- [NPPS16] S. Naderiparizi, A. N. Parks, F. S. Parizi, and J. R. Smith. μ Monitor: In-Situ Energy Monitoring with Microwatt Power Consumption. In *Proceedings of the 2016 IEEE International Conference on RFID (RFID)*, pages 1–8. IEEE, May 2016. doi:10.1109/RFID.2016.7488017.
- [PBLS14] A. Pötsch, A. Berger, C. Leitner, and A. Springer. A Power Measurement System for Accurate Energy Profiling of Embedded Wireless Systems. In *Proceedings of the 2014 IEEE Emerging Technology and Factory Automation (ETFA)*, pages 1–4. IEEE, September 2014. doi:10.1109/ETFA.2014.7005280.
- [PDMJ10] P. Padhy, R. K. Dash, K. Martinez, and N. R. Jennings. A Utility-Based Adaptive Sensing and Multihop Communication Protocol for Wireless Sensor Networks. *ACM Transactions on Sensor Networks*, 6(3):1–39, June 2010. doi:10.1145/1754414.1754423.

- [PGN12] R. Pozza, A. Gluhak, and M. Nati. SmartEye: An Energy-efficient Observer Platform for Internet of Things Testbeds. In *Proceedings of the 7th ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH)*, pages 59–66. ACM, 2012. doi:10.1145/2348688.2348702.
- [PHMM11] C. Peters, J. Handwerker, D. Maurath, and Y. Manoli. A Sub-500 mV Highly Efficient Active Rectifier for Energy Harvesting Applications. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 58(7):1542–1550, July 2011. doi:10.1109/TCSI.2011.2157739.
- [PKHW18] S. Pullwitt, U. Kulau, R. Hartung, and L. C. Wolf. A Feasibility Study on Energy Harvesting from Soil Temperature Differences. In *Proceedings of the 7th International Workshop on Real-World Embedded Wireless Systems and Networks (RealWSN)*, pages 1–6. ACM, November 2018. doi:10.1145/3277883.3277886.
- [PS05] J. A. Paradiso and T. Starner. Energy Scavenging for Mobile and Wireless Electronics. *IEEE Pervasive Computing*, 4(1):18–27, January 2005. doi:10.1109/MPRV.2005.9.
- [RBD13] M. A. Razzaque, C. Bleakley, and S. Dobson. Compression in Wireless Sensor Networks: A Survey and Comparative Evaluation. *ACM Transactions on Sensor Networks*, 10(1):5:1–5:44, November 2013. doi:10.1145/2528948.
- [RKH⁺05] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. Srivastava. Design Considerations for Solar Energy Harvesting Wireless Embedded Systems. In *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 457–462. IEEE, April 2005. doi:10.1109/IPSN.2005.1440973.

-
- [RLSS10] R. R. Rajkumar, I. Lee, L. Sha, and J. Stankovic. Cyber-Physical Systems: The Next Computing Revolution. In *Proceedings of the 47th Design Automation Conference (DAC)*, page 731. ACM, 2010. doi:10.1145/1837274.1837461.
- [RN17] M. I. Rajib and A. Nasipuri. Predictive Retransmissions for Intermittently Connected Sensor Networks with Transmission Diversity. *ACM Transactions on Embedded Computing Systems*, 17(1):1–25, September 2017. doi:10.1145/3092947.
- [Row91] D. Rowe. Applications of Nuclear-Powered Thermoelectric Generators in Space. *Applied Energy*, 40(4):241–271, January 1991. doi:10.1016/0306-2619(91)90020-X.
- [Row06] D. Rowe, editor. *Thermoelectrics Handbook: Macro to Nano*. CRC Press, 1 edition, 2006. doi:10.1201/9781420038903.
- [RSF11] B. Ransford, J. Sorber, and K. Fu. Mementos: System Support for Long-Running Computation on RFID-Scale Devices. In *Proceedings of the 16th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 159–170. ACM, 2011. doi:10.1145/1950365.1950386.
- [RSW19] M. Rottleuthner, T. C. Schmidt, and M. Wählich. Eco: A Hardware-Software Co-Design for In Situ Power Measurement on Low-end IoT Systems. In *Proceedings of the 7rd International Workshop on Energy Harvesting and Energy Neutral Sensing Systems (ENSsys)*, pages 22–28. ACM, 2019. doi:10.1145/3362053.3363495.
- [Sal10] D. Salerno. Ultralow Voltage Energy Harvester Uses Thermoelectric Generator for Battery-Free Wireless Sensors. *LT Journal of Analog Innovation*, 20(3):1–11, 2010.

- [SB93] M. Sugita and W. Brutsaert. Cloud Effect in the Estimation of Instantaneous Downward Longwave Radiation. *Water Resources Research*, 29(3):599–605, March 1993. doi:10.1029/92WR02352.
- [SBC⁺13] J. Sorber, A. Balasubramanian, M. D. Corner, J. R. Ennen, and C. Qualls. Tula: Balancing Energy for Sensing and Communication in a Perpetual Mobile System. *IEEE Transactions on Mobile Computing*, 12(4):804–816, April 2013. doi:10.1109/TMC.2012.52.
- [SC08] S. Silvestre and A. Chouder. Effects of Shadowing on Photovoltaic Module Performance. *Progress in Photovoltaics: Research and Applications*, 16(2):141–149, March 2008. doi:10.1002/pip.780.
- [SCH16] T. Soyata, L. Copeland, and W. Heinzelman. RF Energy Harvesting for Embedded Systems: A Survey of Tradeoffs and Methodology. *IEEE Circuits and Systems Magazine*, 16(1):22–57, December 2016. doi:10.1109/MCAS.2015.2510198.
- [Sci20] Sciencetech Inc. (UHE-45) Large Area Solar Simulator. <http://www.sciencetech-inc.com/uhe-45-large-area-solar-simulator.html>, January 2020.
- [SGT19] L. Sigrist, A. Gomez, and L. Thiele. Tracing Indoor Solar Harvesting. In *Proceedings of the 2nd Workshop on Data Acquisition to Analysis (DATA)*, page 4. ACM, 2019. doi:10.1145/3359427.3361910.
- [SM06] C. M. Sadler and M. Martonosi. Data Compression Algorithms for Energy-Constrained Devices in Delay Tolerant Networks. In *Proceedings of the 4th ACM Conference on Embedded Network Sensor Systems (SenSys)*, pages 265–278. ACM, 2006. doi:10.1145/1182807.1182834.
- [Sny09] G. J. Snyder. Thermoelectric Energy Harvesting. In S. Priya and D. J. Inman, editors, *Energy Harvesting Technologies*, pages 325–336. Springer, Boston, MA, 2009. doi:10.1007/978-0-387-76464-1_11.

-
- [SP09] E. Skoplaki and J. Palyvos. On the Temperature Dependence of Photovoltaic Module Electrical Performance: A Review of Efficiency/Power Correlations. *Solar Energy*, 83(5):614–624, May 2009. doi:10.1016/j.solener.2008.10.008.
- [SPR05] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks. In *Proceedings of the 2005 ACM SIGCOMM Workshop on Delay-Tolerant Networking (WDTN)*, pages 252–259. ACM, 2005. doi:10.1145/1080139.1080143.
- [SR18] G. Shan and B.-H. Roh. Advertisement Interval to Minimize Discovery Time of Whole BLE Advertisers. *IEEE Access*, 6:17817–17825, 2018. doi:10.1109/ACCESS.2018.2817343.
- [Sta14] J. A. Stankovic. Research Directions for the Internet of Things. *IEEE Internet of Things Journal*, 1(1):3–9, February 2014. doi:10.1109/JIOT.2014.2312291.
- [Ste04] J. W. Stevens. Optimal Placement Depth for Air-Ground Heat Transfer Systems. *Applied Thermal Engineering*, 24(2-3):149–157, February 2004. doi:10.1016/j.applthermaleng.2003.09.004.
- [Ste13] J. W. Stevens. Performance Factors for Ground-Air Thermoelectric Power Generators. *Energy Conversion and Management*, 68:114–123, April 2013. doi:10.1016/j.enconman.2012.12.029.
- [STSG16] S. Senni, L. Torres, G. Sassatelli, and A. Gamatie. Non-Volatile Processor Based on MRAM for Ultra-Low-Power IoT Devices. *ACM Journal on Emerging Technologies in Computing Systems*, 13(2), December 2016. doi:10.1145/3001936.
- [SWJ⁺13] M. Salajegheh, Y. Wang, A. A. Jiang, E. Learned-Miller, and K. Fu. Half-Wits: Software Techniques for Low-Voltage Probabilistic Storage on Microcontrollers with NOR Flash Memory. *ACM Transactions*

- on Embedded Computing Systems*, 12(2s):1–25, May 2013. doi:10.1145/2465787.2465793.
- [Tex13] Texas Instruments. Bq25570: Nano Power Boost Charger and Buck Converter for Energy Harvester Powered Applications, 2013.
- [TP11] Y. K. Tan and S. K. Panda. Energy Harvesting From Hybrid Indoor Ambient Light and Thermal Energy Sources for Enhanced Performance of Wireless Sensor Nodes. *IEEE Transactions on Industrial Electronics*, 58(9):4424–4435, September 2011. doi:10.1109/TIE.2010.2102321.
- [TSM⁺17] M. Thielen, L. Sigrist, M. Magno, C. Hierold, and L. Benini. Human Body Heat for Powering Wearable Devices: From Thermal Energy to Application. *Energy Conversion and Management*, 131:44–54, January 2017. doi:10.1016/j.enconman.2016.11.005.
- [TWK⁺14] T. T. Toh, S. W. Wright, M. E. Kiziroglou, P. D. Mitcheson, and E. M. Yeatman. A Dual Polarity, Cold-Starting Interface Circuit for Heat Storage Energy Harvesters. *Sensors and Actuators A: Physical*, 211:38–44, May 2014. doi:10.1016/j.sna.2014.02.037.
- [UMJ12] S. Upadhyay, S. Mishra, and A. Joshi. A Wide Bandwidth Electronic Load. *IEEE Transactions on Industrial Electronics*, 59(2):733–739, February 2012. doi:10.1109/TIE.2011.2148680.
- [UYE⁺15] S. Ulukus, A. Yener, E. Erkip, O. Simeone, M. Zorzi, P. Grover, and K. Huang. Energy Harvesting Wireless Communications: A Review of Recent Advances. *IEEE Journal on Selected Areas in Communications*, 33(3):360–381, March 2015. doi:10.1109/JSAC.2015.2391531.
- [VGB07] C. M. Vigorito, D. Ganesan, and A. G. Barto. Adaptive Control of Duty Cycling in Energy-Harvesting Wireless Sensor Networks. In *Proceedings of the*

-
- 4th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, pages 21–30. IEEE, June 2007. doi:10.1109/SAHCN.2007.4292814.
- [VS19] G. Verma and V. Sharma. A Novel Thermoelectric Energy Harvester for Wireless Sensor Network Application. *IEEE Transactions on Industrial Electronics*, 66(5):1–1, May 2019. doi:10.1109/TIE.2018.2863190.
- [WCK⁺14] C. Wang, N. Chang, Y. Kim, S. Park, Y. Liu, H. G. Lee, R. Luo, and H. Yang. Storage-Less and Converter-Less Maximum Power Point Tracking of Photovoltaic Cells for a Nonvolatile Microprocessor. In *Proceedings of the 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 379–384. IEEE, January 2014. doi:10.1109/ASPDAC.2014.6742919.
- [WCW⁺13] W. Wang, V. Cionca, N. Wang, M. Hayes, B. O’Flynn, and C. O’Mathuna. Thermoelectric Energy Harvesting for Building Energy Management Wireless Sensor Networks. *International Journal of Distributed Sensor Networks*, 9(6):232438, June 2013. doi:10.1155/2013/232438.
- [WD12] S. A. Whalen and R. C. Dykhuizen. Thermoelectric Energy Harvesting from Diurnal Heat Flow in the Upper Soil Layer. *Energy Conversion and Management*, 64:397–402, December 2012. doi:10.1016/j.enconman.2012.06.015.
- [WLW⁺16] Y. Wang, Y. Liu, C. Wang, Z. Li, X. Sheng, H. G. Lee, N. Chang, and H. Yang. Storage-Less and Converter-Less Photovoltaic Energy Harvesting with Maximum Power Point Tracking for Internet of Things. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 35(2):173–186, February 2016. doi:10.1109/TCAD.2015.2446937.
- [WMM⁺13] A. S. Weddell, M. Magno, G. V. Merrett, D. Brunelli, B. M. Al-Hashimi, and L. Benini. A Survey of Multi-Source Energy Harvesting Systems. In *Proceedings*

- of the 2013 Design, Automation and Test in Europe Conference and Exhibition (DATE), pages 905–908. IEEE, 2013. doi:10.7873/DATE.2013.190.
- [WTZE09] R. Wohlers, N. Trigoni, R. Zhang, and S. Ellwood. TwinRoute: Energy-Efficient Data Collection in Fixed Sensor Networks with Mobile Sinks. In *Proceedings of the 10th International Conference on Mobile Data Management: Systems, Services and Middleware*, pages 192–201. IEEE, 2009. doi:10.1109/MDM.2009.30.
- [Yak11] A. Yakovlev. Energy-Modulated Computing. In *Proceedings of the 2011 Design, Automation and Test in Europe Conference and Exhibition (DATE)*, pages 1–6. IEEE, March 2011. doi:10.1109/DATE.2011.5763216.
- [YLYC18] J. Yan, X. Liao, D. Yan, and Y. Chen. Review of Micro Thermoelectric Generator. *Journal of Microelectromechanical Systems*, 27(1):1–18, February 2018. doi:10.1109/JMEMS.2017.2782748.
- [YZLM14] S. Yu, B. Zhang, C. Li, and H. Mouftah. Routing Protocols for Wireless Sensor Networks with Mobile Sinks: A Survey. *IEEE Communications Magazine*, 52(7):150–157, July 2014. doi:10.1109/MCOM.2014.6852097.
- [ZGL13] P. Zhang, D. Ganesan, and B. Lu. QuarkOS: Pushing the Operating Limits of Micro-Powered Sensors. In *Proceedings of the 14th USENIX Conference on Hot Topics in Operating Systems*. USENIX, 2013.
- [ZLT16] J. Zhang, Z. Li, and S. Tang. Value of Information Aware Opportunistic Duty Cycling in Solar Harvesting Sensor Networks. *IEEE Transactions on Industrial Informatics*, 12(1):348–360, February 2016. doi:10.1109/TII.2015.2508745.
- [ZX13] R. Zhou and G. Xing. Nemo: A High-Fidelity Non-invasive Power Meter System for Wireless Sensor Networks. In *Proceedings of the 12th ACM/IEEE International Conference on Information Processing in*

Sensor Networks (IPSN), pages 141–152, April 2013.
doi:10.1145/2461381.2461401.

List of Publications

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

Gomez A, Sigrist L, Magno M, Benini L, Thiele L. **Dynamic Energy Burst Scaling for Transiently Powered Systems.** *In Proceedings of the 2016 Design, Automation and Test in Europe Conference and Exhibition (DATE)*. IEEE. 2016. (Chapter 5)

Sigrist L, Gomez A, Lim R, Lippuner S, Leubin M, Thiele L. **RocketLogger: Mobile Power Logger for Prototyping IoT Devices: Demo Abstract.** *In Proceedings of the 14th International Conference on Embedded Network Sensor Systems (SenSys)*. ACM. 2016. (Chapter 2)

Gomez A, Sigrist L, Schalch T, Benini L, Thiele L. **Efficient, Long-Term Logging of Rich Data Sensors using Transient Sensor Nodes.** *ACM Transactions on Embedded Computing Systems*. ACM. 2017. (Chapter 5)

Sigrist L, Gomez A, Lim R, Lippuner S, Leubin M, Thiele L. **Measurement and Validation of Energy Harvesting IoT Devices.** *In Proceedings of the 2017 Design, Automation and Test in Europe Conference and Exhibition (DATE)*. IEEE. 2017. (Chapter 2)

Sigrist L, Thiele L. **Design Support for Energy Harvesting Driven IoT Devices.** *In Proceedings of the IDEA League Doctoral School on Transiently Powered Computing*. Achener Informatik Berichte. 2017. (Chapter 3)

Sigrist L, Stricker N, Bernath D, Beutel J, Thiele L. **Thermoelectric Energy Harvesting from Gradients in the Earth Surface.** *IEEE Transactions on Industrial Electronics*. IEEE, 2019. (Chapter 4)

Sigrist L, Ahmed R, Gomez A, Thiele L. **Harvesting-Aware Optimal Communication Scheme for Infrastructure-Less Sensing.** *ACM Transaction on Internet of Things*. [under review] (Chapter 6)

Sigrist L, Gomez A, Leubin M, Beutel J, Thiele L. **Environment and Application Emulation Testbed for Reliable Energy-Harvesting System Design.** *IEEE Transactions on Industrial Electronics*. [to be submitted] (Chapter 3)

The following list includes publications that are not part of this thesis.

Sigrist L, Giannopoulou G, Huang P, Gomez A, Thiele L. **Mixed-criticality Runtime Mechanisms and Evaluation on Multicores.** *In Proceedings of the 2015 Real-Time and Embedded Technology and Applications Symposium (RTAS).* IEEE. 2015.

Nehani J, Brunelli M, Magno M, Sigrist L, Benini L. **An Energy Neutral Wearable Camera with EPD Display.** *In Proceedings of the 2015 Workshop on Wearable Systems and Applications (WearSys).* ACM. 2015.

Magno M, Brunelli D, Sigrist L, Andri R, Cavigelli L, Gomez A, Benini L. **InfiniTime: Multi-Sensor Wearable Bracelet with Human Body Harvesting.** *Sustainable Computing: Informatics and Systems.* Elsevier. 2016.

Thielen, M, Sigrist L, Magno M, Hierold C, Benini L. **Human Body Heat for Powering Wearable Devices: From Thermal Energy to Application.** *Energy Conversion and Management.* Elsevier. 2017.

Gomez A, Sigrist L, Schalch T, Benini L, Thiele L. **Wearable, Energy-Opportunistic Vision Sensing for Walking Speed Estimation.** *In Proceedings of the 2017 IEEE Sensors Applications Symposium (SAS).* IEEE. 2017.

Sigrist L, Gomez A, Lim R, Lippuner S, Leubin M, Thiele L. **In-situ-Validierung von Energy-Harvesting.** *DESIGN&ELEKTRONIK.* WEKA Fachmedien. 2017.

Magno M, Sigrist L, Gomez A, Cavigelli L, Libri A, Popovici E, Benini L. **SmarTEG: An Autonomous Wireless Sensor Node for High Accuracy Accelerometer-Based Monitoring.** *Sensors.* MDPI. 2019.

Ahmed R, Buchli B, Draskovic S, Sigrist L, Kumar P, Thiele L. **Optimal Power Management with Guaranteed Minimum Energy Utilization for Solar Energy Harvesting Systems.** *ACM Transactions on Embedded Computing Systems.* ACM. 2019.

Sigrist L, Gomez A, Thiele L. **Tracing Indoor Solar Harvesting.** *In Proceedings of the 2nd Workshop on Data Acquisition to Analysis (DATA).* ACM, 2019.

Curriculum Vitae

