Wireless Low Power Edge Devices and Transient Computing Systems with Energy Harvesting

Drahtlose Low Power Endgeräte und Transient Computing Systeme mit Energy Harvesting

Vom Fachbereich Elektrotechnik und Informationstechnik
der Rheinland-Pfälzischen Technischen Universität Kaiserslautern-Landau
zur Verleihung des akademischen Grades
Doktor der Ingenieurwissenschaften (Dr.-Ing.)
genehmigte Dissertation

von

Carl Christian Rheinländer

geboren in Simmern

Tag der Einreichung:	28. Mai 2024
Tag der mündlichen Prüfung:	07. März 2025
Dekan des Fachbereichs:	Prof. Dr.-Ing. Daniel Görges
Vorsitzender der Prüfungskommission:	Prof. Dr.-Ing. Daniel Görges
1. Berichterstatter:	Prof. Dr.-Ing. Norbert Wehn
2. Berichterstatter:	Prof. Dr. rer. nat. Wilhelm Stork

DE-386

Eidesstattliche Erklärung

Ich erkläre an Eides statt,

dass ich die vorliegende Dissertation selbst angefertigt habe und alle von mir benutzten Hilfsmittel angegeben habe,

dass ich die Dissertation oder Teile hiervon noch nicht als Prüfungsarbeit für eine staatliche oder wissenschaftliche Prüfung eingereicht habe, und

dass ich nicht die gleiche oder eine andere Abhandlung bei einem anderen Fachbereich oder einer anderen Universität als Dissertation eingereicht habe.

Kaiserslautern, 28. Mai 2024

Carl Christian Rheinländer

Abstract

Since the prediction of Gordon Moore in 1965, known as Moore's Law [1], the integration density and performance of integrated circuits have increased continuously. As a result, integrated circuits are enabling ever-new possibilities for decades and are being used in more and more areas of applications. This led to a fast growing demand of microelectronic systems ranging from, e.g., edge devices for the Internet of Things (IoT) over industrial control systems up to computationally intensive machine learning applications.

However, within the last years, the increased environmental awareness and the efforts to create a more sustainable world have shed light on the darker side of this development: The increasing energy demand for computing. In fact, the global demand for computational energy and battery capacity is increasing exponentially. In order to counteract this development, experts call for fundamental changes in microelectronic systems design that enable significantly better energy efficiencies of electronic devices [2]. Suggestions range from an increased use of advanced computing platforms that allow for application-specific hardware-software co-designs up to an improvement of edge computing capabilities to gradually replace conventional computing. More radical proposals request for a consistent utilization of renewable energy sources [3] or propose specific paradigm changes in commonly battery-powered application fields by calling for a battery-less internet of things [4].

This thesis proposes new design approaches for energy and environmentally friendly embedded systems in different areas of application that are fit for the future. For the industrial area, this thesis addresses model predictive control (MPC) systems. A new optimization strategy for the resource-intensive MPC algorithms is demonstrated to enable their implementation on low power and low-cost FPGAs. Regarding edge computing, new design approaches are presented to improve the computing capabilities for mobile battery-powered devices. These approaches enable ultra-low power communication security features for IoT systems and present new precise time synchronization principles for wireless sensor networks.

A clear emphasis of this thesis is on the improvement of embedded systems that are solely powered by energy harvesting. Such battery-less systems are also known as intermittent or transient computing systems, as they do not have a reliable energy basis. One of the major issues that hamper the widespread use of this sustainable system concept is the *reactive* nature of state-of-the-art transient computing approaches, i.e., the system is completely and utterly dependent on the transiency of its energy harvesting source. As a solution, this thesis proposes new hardware and software design methods towards a *proactive* transient computing. These methods are dedicated to kinetic and thermal energy harvesting technologies and enable a predictive and adaptive power loss management to pave the road for yet infeasible applications fields.

Acknowledgments

Many friends, colleagues and students have in one way or another contributed to the creation of this work. On behalf of all, I would like to thank some individuals explicitly.

Especially I would like to thank my doctoral adviser Prof. Dr.-Ing. Norbert Wehn for the opportunity to work in the interesting field of transient computing systems and wireless edge devices in an inspiring and congenial work atmosphere as well as for all the enlightening conversations and productive discussions. Many thanks to Prof. Dr. rer. nat. Wilhelm Stork for his interest in my thesis and for serving on my thesis committee and to Prof. Dr.-Ing. Daniel Görges for chairing this committee.

I am very grateful for my friends and colleagues I had the chance to work with in the Microelectronic Systems Design Research Group: Dr.-Ing. Imran Ali, Dr.-Ing. Matthias Alles, Dr.-Ing. Dipl.-Phys. Christian Brugger, M.Sc. André Lucas Chinazzo, Dr.-Ing. Christian De Schryver, M.Sc. Johannes Feldmann, M.Sc. Muhammad Mohsin Ghaffar, M.Sc. Oliver Griebel, Dipl.-Ing. Matthias Herrmann, Dr.-Ing. Thomas Ilnseher, M.Eng. Lucas Johannsen, Prof. Dr.-Ing. Matthias Jung (DL9MJ), Dipl.-Inf. Claus Kestel, PD Dr.-Ing. habil. Frank Kienle, M.Sc. Kira Kraft, Dr.-Ing. Lukas Krupp, M.Sc. Jan Lappas, M.Sc. Frederik Lauer, Dr.-Ing. Timo Lehnigk-Emden, Dipl.-Ing. Dominik Loroch, Dr.-Ing. Deepak M. Mathew, M.Sc. Jonas Ney, Dr.-Ing. Vladimir Rybalkin, Dr.-Ing. Mohammadsadegh Sadri, Dr.-Ing. Philipp Schläfer, M.Sc. Maximilian Schöffel, Dr.-Ing. Stefan Scholl (DC9ST), M.Sc. Lukas Steiner, Dr.-Ing. Chirag Sudarshan, M.Sc. Menbere Tekleyohannes, Dr.-Ing. Javier Alejandro Varela, Dipl.-Math. Uwe Wasenmüller, Dr.-Ing. Christian Weis, Dr.-Ing. Stefan Weithoffer, Dipl.-Ing. Sebastian Wille, Dr.-Ing. Renhai Xiong and Eng. Éder Ferreira Zulian.

I would like to express my deepest gratitude to our secretary Martina Jahn for her constant helpfulness, her assistance in every critical situation and her dedicated support on all conceivable issues. I would also like to thank Dipl.-Ing. (FH) Hans-Peter Goldhammer, Dipl.-Ing. (FH) Roland Volk and Andreas Christmann for supporting me in all kinds of technical issues with their comprehensive expertise.

Special thanks to all the external scientists I worked with, especially to Dr. Marcus Schmidt and Prof. Dr. Thomas Jaitner from the TU Dortmund, Dr.-Ing. Felix Berkel from the Chair of Control Systems, Dipl.-Ing. Yvo Stiemcke and Jun. Prof. Dr.-Ing. Stefan Thielen from the Department of Mechanical and Process Engineering, as well as to all the supporting students, especially to Maycon Douglas, Vitor Kieling and Martin Schultheis.

Also, many thanks to the students of the AktionsBündnis and the professors at the Hochschule Kaiserslautern for their great support, especially Prof. Dr.-Ing. Günter Biehl.

I am deeply grateful to my grandfathers Werner Connertz and Karl Rheinländer, who awakened my enthusiasm for technical relationships in electronics in my childhood, to my parents Annemarie and Carl Christian Rheinländer, who enable me to experiment intensively in this field using every conceivable means, and to my wife Ida Rheinländer, who always encouraged and supported me and had my back even in the most difficult times.

It was only with all their support that I was able to write this thesis.

Contents

Chapter 1

Introduction

> ”Man needs difficulties, they are necessary for health.”
>
> Carl Gustav Jung

For the last decades, microelectronic systems have grown rapidly in terms of computing performance and integration complexity. Microelectronics predominantly depend on the capabilities of *integrated circuits* (ICs). ICs range from programmable general-purpose *Central Processing Units* (CPUs) and *Microcontroller Units* (MCUs), over programmable controllers for specific purposes such as *Graphics Processing Units* (GPUs) and *Digital Signal Processors* (DSPs) as well as *application-specific integrated circuits* (ASICs) up to programmable logic devices like *Field Programmable Gate Arrays* (FPGAs). Modern heterogeneous devices host different types of aforementioned systems and memories together, which are interconnected on the chip, forming *System on Chips* (SoCs) and *Network on Chips* (NoCs). All these different types of ICs share the same and comparably trivial base component: The transistor. Multiple transistors are combined to form logic gates, which are the generic basis of any digital circuitry. Generally, the functionality and performance of an IC scales with the number of incorporated transistors. For many decades, the development of ICs can be seen as a straight and continuous improvement towards more incorporated transistors and improved computational capability. This trend was quantitatively predicted by Gordon Moore in 1965 in the way that the number of transistors on a chip will double every approximately two years and therewith the performance of the chips as well [1]. At the same time, transistor sizes have shrunken continuously such that the integration density was increasing. Concurrently, the manufacturing cost per transistor has massively decreased. Therefore, the computing performance of ICs has increased over the years at nearly no increase in cost. However, this trend has ended as physical limits have been reached that led to advanced technologies and fabrication process requirements.

1.1 Challenges and Motivation

The transistor size reduction trend of the conventional planar CMOS (complementary metal-oxide semiconductor) fabrication process has significantly slowed down as the physical dimensions of transistors came closer to the silicon atom level. When the size went below the 20nm range, the transistors suffered from increased leakage and quantum tunneling effects that, e.g., led to unreliable off-states. The invention of the Fin-FET helped to counteract these effects to continue the decreasing of the size further. Recent developments however have shown, that the Fin-FET era is ending and, in order to realize transistor sizes below 7nm, the transistor geometry was changed again and the GAA-FET (Gate-all-around-FET) technology has been released.

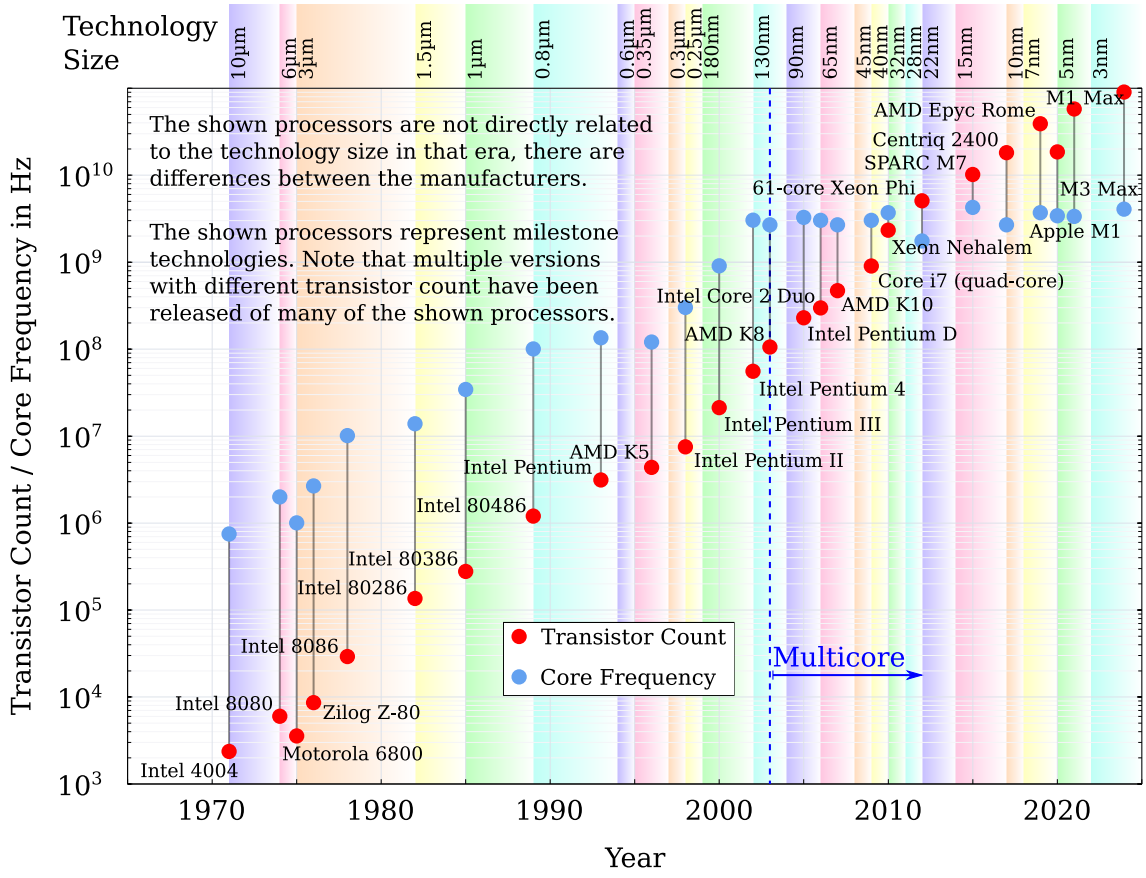

Figure 1.1: Time line of the microprocessor evolution with chosen devices with respect to transistor count and technology size (data obtained from [2,5–8], Oracle, Intel, AMD, Apple and Qualcomm).

In 1974, Dennard et al. postulated a scaling law that describes the potentials of shrinking transistor sizes [9]. It describes how smaller sizes enable a lowering of the operating voltage and thus an increasing of the operating frequency without raising the power density. Therefore, the performance of commercial singe-core CPUs has increased rapidly for a few decades. However, about 30 years later, this trend stopped as the operating voltage could not be decreased any further. As a result, the power density increased with the frequency

and led to critical levels of heat dissipation that more and more hit the limits of cooling methods. Therefore, increasing the performance by increasing the operating frequency reached its limit where the power density became critical (*Power Wall*).

Nowadays, the power density is a key constraint in IC design. To continue a growth in performance, computational operations have been distributed on the die surface, by incorporating multiple processing cores that run at a lower operating voltage. This led to an increase in parallelism, which enabled to increase the overall performance without increasing the clock frequency. Figure 1.1 shows how this multicore approach affected the evolution of microprocessor architectures over the last decades.

However, the performance improvement due to parallel computing, obtained, e.g., from multicore architectures, is limited. According to *Amdahl's Law*, an increase in the number of cores and therefore parallel computational capability, does not implicate the same level of increase in performance. A further development was to incorporate dedicated hardware subsystems, e.g., accelerators, to which specific operations are outsourced. The development and deployment of such integrable hardware subsystems and circuits, known as *intellectual property* (IP), has even become a core business model for many hardware companies meanwhile. With this strategy, the integration density can be increased beyond the critical power density.

All these historic and latest developments have achieved a continuous improvement of the energy performance of computing — physically expressible as Joule per Bit $\frac{J}{Bit}$ — despite the discussed limits.

Growing Demand for Computing Capacity

The enormous progress of integrated circuits and the continuous reduction of cost for electronic components led to a massive increase in deployment of microelectronics for wearable devices, smart home, cars and consumer electronics [10]. The microcontroller as one of the common core components is found in vehicles, cellphones and smartwatches, dishwashers and washing machines, computer peripherals, smart home systems, *Internet of Things* (IoT) devices, medical appliances and life support systems. Most microcontrollers shipped are employed in the embedded sector, those in *Personal Computers* (PCs) only made up around 6% in the early 2010s [11] and fell to 2% in 2018 [12]. The total number of microcontrollers is not easy to assess, but semiconductor companies state that the average consumer gets in touch of more than 100 products that employ an MCU every day [13]. This number is growing as more and more electronic products are deployed every day [14, 15]. Over 28 billion units of MCUs have been shipped globally in the year 2020 [16]. In 2022, the application area of consumer electronics and telecommunication dominated the microcontroller market with a market share of almost 33% [17]. Figure 1.2 shows how the U.S. MCU market is expected to grow until 2027 for the different application areas.

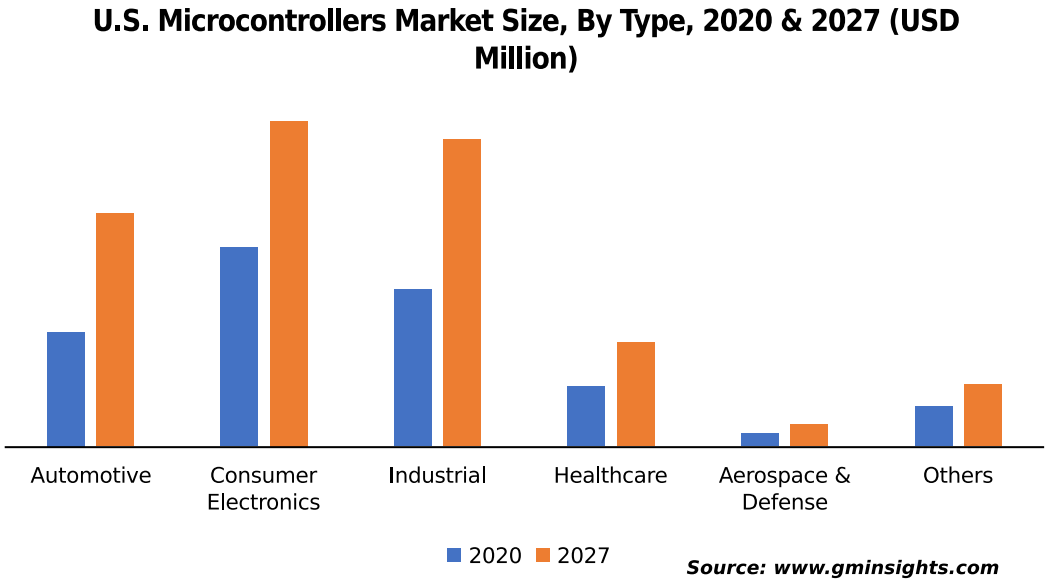

Figure 1.2: Relative MCUs market size in the U.S. for 2020 and 2027 [18].

The future *Compound Annual Growth Rates* (CAGRs) are reported to be slightly below 10% from 2021 to 2026 [14, 19], over around 10% from 2021 to 2028 [16, 20] up to 14% from 2021 to 2031 [15]. Respective MCU market values are calculated to have been around USD 20 billion in 2021 [15, 16] and are expected to reach around USD 36 billion [21] to over USD 48 billion by 2027 [22] up to over USD 74 billion by 2031 [15]. The main driver for the huge growth is the increasing demand in the automotive sector, where a growing demand exists for driver assistance and safety systems due to growing standards, autonomous driving and digitalization of automobiles [15, 22, 23]. This sector is followed by consumer and telecommunications electronics, incorporating all kinds of IoT devices. The number of deployed IoT devices was around 8.6 billion in 2019 and is expected to grow to about 30 billion in 2030 [24]. In the industrial sector, main drivers are new demands for flexibility and agility [23], emerging industrial interconnection known as *Industrial Internet of Things* (IIoT) infrastructures, as well as all further technical approaches known as *Industry 4.0*.

Energy-Related Consequences

The discussed trends entail an ever-growing global demand for computing, measured in *zetta instructions per second* (ZIPS)[1] as shown in Figure 1.3.

[1] ZIPS: 10^{21} compute instructions per second

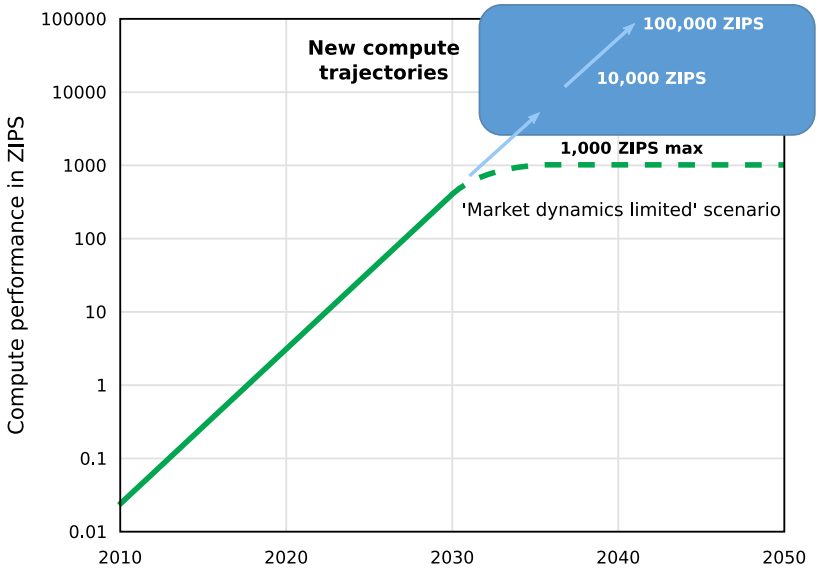

Figure 1.3: Global capacity to compute information in ZIPS[2](graphic based on [2] (page 123)).

Up to now, the energy demand for this computation has grown significantly faster — approximately doubling every three years — than the global energy production capacity [2,3]. This development is expected to continue despite the substantial progress in chip-level technology during the last decades, i.e., despite the continuous improvement of the energy performance of computing [2, 3, 25]. According to the *Decadal Plan for Semiconductors* [2] and shown in Figure 1.3 and Figure 1.4, this will lead to a *market-dynamics-limited* scenario in the near future, unless fundamental changes are achieved.

These changes need to radically improve the energy efficiency of microelectronic systems and change the compute trajectories. Otherwise, there will be a significant increase in global energy production to serve the energy demand for computing in the future. This demand already constitutes more than 2% of the global energy demand [3].

Besides the need for an increased production of electric grid energy, the massive increase in produced mobile microelectronic devices entails a fast-growing demand in mobile electrical energy sources like batteries. This is clearly illustrated by the global market trends for the most popular battery technology: The *Lithium-Ion battery* (LIB). First commercially introduced by Sony in 1991, this technology is now more or less the state-of-the-art battery base technology [26]. The growing demand of electric vehicles is a significant contributor to an extreme growth in battery capacity demand [26, 27], others are stationary energy storage and consumer products [28]. Over the last ten years, the global demand has grown from $0.5GWh$ in 2010 to over $500GWh$ in 2020 [28]. Scientific studies in [29] and [30] expect the capacity demand to approximately double every five years, and market researchers expect it to grow eleven-fold from 2020 to 2030 [26]. The LIB production capacity is expected to grow to almost $3000GWh$ by 2030 [27]. In the consumer sector

[2]ZIPS: 10^{21} compute instructions per second

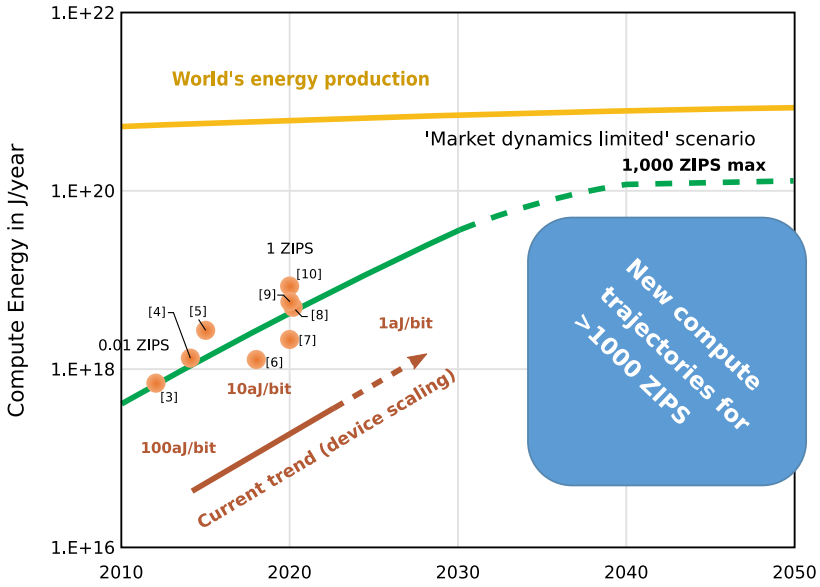

Figure 1.4: Despite the improving energy performance $\frac{J}{Bit}$, the current trend of demanded compute energy (green line) grows way faster than the world's energy production (yellow line) (graphic based on [2] (page 123), references in the graph relate to the references in [2]).

alone, the production capacity is expected to grow from $60GWh$ in 2020 to $372GWh$ in 2030 (Figure 1.5).

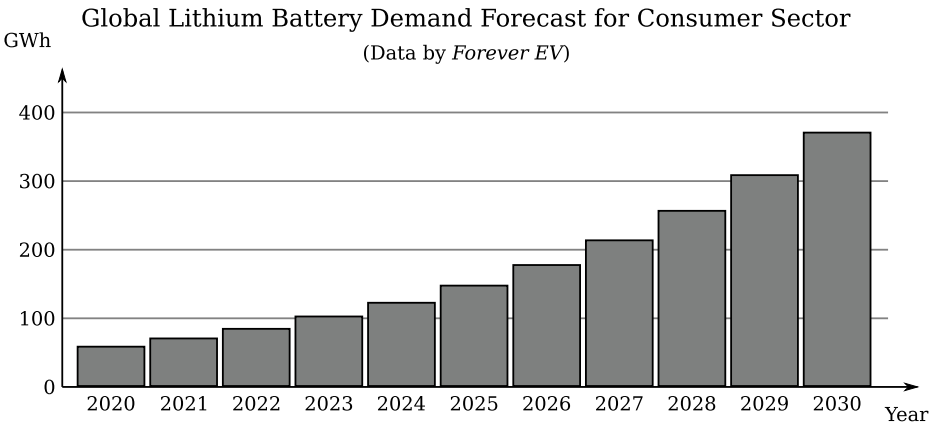

Figure 1.5: Forecast of global Lithium battery demand for consumer sector [27].

This results in an enormous increase in demand for resources like Lithium, which is accompanied by myriads of adverse effects on the environment like destruction of natural landscape, poisoning of ground water, consumption of fossil energy carriers, environmental pollution and exorbitant emission of carbon dioxide.

Energy-Specific Challenges

This globally growing demand for resources and energy for computation is gaining more and more attention in the context of the efforts to create a more sustainable and environmentally friendly world. In order to prevent a dramatically increase of the world's energy production, radically new computing paradigms must be discovered that provide extremely increased energy efficiencies of microelectronic systems [2]. It is now, among others, up to the experts of microelectronic systems and the technology to come up with new approaches for the future of microelectronic devices. For instance, by making use of more application-specific and optimized co-design of embedded hardware and software, or by shifting conventional computing more and more towards edge devices (edge computing), up to an increased utilization of renewable energy sources [3]. To repress the fast-growing demand for batteries like LIB in the consumer sector, the utilization of renewable energy sources has to gain solid ground in this fast growing market to foster the path for mobile battery-less electronic devices [4].

This thesis contributes to this by proposing new solutions for embedded microelectronic systems on the system design level. The proposed solutions cover different application areas and point out specific optimization strategies and improvements for respective embedded systems:

- New approaches for *proactive* transient computing on *Power-Neutral* energy harvesting systems. *Power-Neutral* energy harvesting systems can be seen as the most extreme approach of energy efficiency improvements since these systems require no conventional power supply at all. Instead, an energy harvester scavenges the entire operating energy from its environment. As *Power-Neutral* systems do not buffer the harvested energy, these systems don't even employ energy buffers like batteries and therefore represent a pioneering future design principle towards battery-less microelectronic systems.
- Enhancing the battery lifetime of mobile sensor devices and IoT edge devices by dedicated hardware and software design strategies. Therewith improving the potential and capability of ultra-low power edge computing.
- Improving the energy efficiency by application-specific hardware implementation strategies for computationally intensive operations, enabling a low power and embedded hardware-based implementation of a model predictive control system for infrastructure facilities instead of a conventional, energy-intensive software-based implementation on an industrial computer.

1.2 New Contributions and Thesis Structure

Categorization of Embedded Systems

Energy-related scientific challenges of embedded microelectronic systems strongly depend on the chosen energy or power source, which can be categorized into three basic options: Energy harvesting, battery and energy grid. This choice depends on the system requirements and intended use case, and determines the boundary conditions and constrains for the system design as shown in Table 1.1.

Power Source	**Energy Harvesting**	**Battery**	**Energy Grid**
Available Power Budget	Low	Middle	High
Power Supply Reliability	Transient	Limited	Unlimited
Energy-Specific Design Strategy	Maximal exploitation of available energy	Optimize for long battery lifetime	Optimize for best energy efficiency

Table 1.1: The different power source options and their related design constrains and energy-specific design strategy.

The field *Energy-Specific Design Strategy* in Table 1.1 describes the design strategy to achieve the best overall system energy efficiency. This overall system energy efficiency incorporates the energy efficiency of the load system together with the utilization efficiency of the power source. It represents the optimum of the combination of both efficiencies that can be achieved under the given constrains. These constrains are fundamentally different for the mentioned power source options:

- Energy harvesting-powered systems: Even though the energy that can be harvested is theoretically unlimited, the available power budget is typically very low for embedded systems. Since its availability is determined by the environmental conditions, neither the power level nor its time of availability are predictable. The absence of an energy buffer furthermore means that all harvested power is transient and only available instantly. To achieve the best overall system efficiency, all incoming energy must be exploited as efficient as possible. Therefore, the system must be prepared for all possible situations at all time, e.g., from sudden high-power income levels down to sudden power losses, and adopt quickly to changing conditions.
- Battery-powered systems: The system design is restricted with respect to the chosen battery technology, chemistry and size. Different technologies come with different characteristics like output voltage level, internal impedance and usable capacity, i.e., the total available amount of energy. Most of these characteristics vary over time and depend on the state of charge, temperature and current load. For example, many

battery chemistries show a degradation of capacity or increase in internal impedance at high discharge rates. Thus, the total available amount of energy depends on the maximum current load as well as the duty cycle of the load system. This limits the design space of the load system, especially regarding its energy efficiency strategy, as the specific power source characteristics and resulting interdependencies between source and load must be considered carefully. In general, an overall energy efficient system therefore is a compromise between power source and load system and can be expressed as the optimal interaction between an energy efficient load system design and an energy-efficient operation of the power source, which leads to the longest possible battery lifetime. Therefore, in order to express the overall system energy efficiency, the total battery lifetime is the more meaningful metric than the plain energy efficiency measures of the load system.

- Energy grid-powered systems can often be optimized for best energy efficiency without taking specific characteristics of the power source into account. Regarding embedded systems, the energy grid can usually be seen as a low-impedance source that shows no interaction with the current load of the embedded system. This allows a broad flexibility in terms of system design, e.g., the choice of hardware platform, as the power source has almost no limitations in terms of output voltage and current load.[3]

Based on these three types of power supply, this work defines three categories of embedded systems that are addressed separately in this thesis as shown in Table 1.2.

Power Source	Energy Harvesting	Battery	Energy Grid
Category of Embedded System	**Autonomous Embedded Systems**	**Mobile Embedded Systems**	**Stationary Embedded Systems**
Power Supply Characteristics	Powered by energy harvesting without energy buffers	Powered by primary or rechargeable battery cells	Powered by the energy grid
Addressed in	Chapter 2	Chapter 3 (rechargeable battery) Chapter 4 (primary battery)	Chapter 5

Table 1.2: The different categories of embedded systems addressed in this work.

Addressed Scientific Challenges

The design constrains of the discussed categories of embedded systems pose particular scientific system design challenges that have been addressed in this work:

[3] In this argumentation, the design of dedicated *alternating current* (AC)-adapters or other voltage converters to electrically connect stationary embedded systems to the energy grid is considered as a sub component in the overall system design and thus also subject to energy-efficient optimization.

- **Autonomous Embedded Systems** (addressed in Chapter 2)

  Battery-free and energy harvesting-powered microelectronic systems go beyond the efforts of ultra-low power system design, as they scavenge their entire operational energy from the environment instead of relying on a conventional energy reservoir. However, due to the transient nature of energy harvesting, these systems do not have a continuous power basis. To hide this transiency to the load system and thus to simplify the system design, the *Energy-Neutral* design concept utilizes energy storage units like capacitors to buffer the varying power input over time. In contrast, the *Power-Neutral* design concept waives any energy buffers to mitigate the influence of buffer-related drawbacks such as aging effects as well as capacity and leakage losses due to the continuous charge and discharge cycles. In return, the load system becomes a transiently powered system that must be designed to weather unforeseen power losses all the time. This leads to design challenges, which are significantly different to those in battery-powered systems:

  - Ensuring forward progress of the application despite sudden power losses, and guaranteeing the finalization of atomic operations on an unreliable energy basis.
  - Typically, low power outputs and short output durations of energy harvesters define strict power and processing time limits, which usually are not calculable or predictable.
  - Frequent resets and restarts exacerbate the negative influence of cold-start overheads on the energy efficiency of the system. This particularly takes effect in the design of transiently powered systems out of off-the-shelf components. These components are usually not intended for unforeseen power supply interruptions nor optimized for low cold-starts as their low power strategy is to rest in sleep modes most of the time rather than in total power off conditions.
  - Different types of energy harvesters have significantly different characteristics in terms of power output intensity and duration. Furthermore, the retroactive impacts of the system load on the harvester and its physical excitation are very specific to the transduction, e.g., breaking effects of electromagnetic transducers or thermal conductivity alteration of thermoelectric transducers (Thomson effect). Therefore, sophisticated system designs require harvester-specific optimizations and holistic analyses.

- **Mobile Embedded Systems** (addressed in Chapter 3 and Chapter 4)

  The progress of battery technologies regarding their energy density did not even partially keep up with the progress of microelectronics regarding their integration density and performance enhancement. As a result, the battery nowadays tends to take up a significant proportion of space inside electronic edge devices such as wearable sensor appliances and mobile IoT devices. Together with further requirements like

use case-dependent dimensional constrains, the battery becomes the main limiting factor in mobile electronic systems. In other words, the limited energy budget of the battery on the one hand and the desired high system performance on the other hand constitute the major design challenges for future edge computing:

- To provide the maximal performance on a strictly limited energy budget, application-specific ultra-low power design methods for hardware and software are required. In wearable *wireless sensor networks* (WSNs), e.g., this involves a specific selection of components, a space-constrained construction and an appropriate software design to enable precise long-term sensing of movement data despite the limited battery capacity.
- In the IoT, the number of interconnected battery-powered edge devices has grown continuously within the last years. However, such energy-constrained devices often lack of sophisticated security mechanisms due to the computational complexity of cryptographic algorithms and the related energy demand. Therefore, insufficient communication security has been valued as a serious risk for IoT edge computing and as one of the most critical challenges in this field.
- As the limited energy budget also limits the wireless connectivity, appropriate low power wireless technologies like *Bluetooth Low Energy* (BLE) come with low bandwidth and restricted flexibility. The access to the packet transfer timing schedule of BLE is blocked to ensure a reliable connectivity, which however makes precise wireless synchronization extremely difficult for wearable WSNs. Furthermore, the bandwidth limitation complicates the implementation of *Transport Layer Security* (TLS)-based security mechanisms that rely on mutual authentication and thus the timely transfer of long certificates and handshake messages.

- **Stationary Embedded Systems** (addressed in Chapter 5)

  Stationary embedded systems are usually a sub component of other systems. In the consumer sector they often constitute the control unit of electrical appliances, whereas in the industrial sector they take over control tasks of processing plants and other industrial processes. Many industrial control tasks place high requirements on the control system like real-time capability, guaranteed control stability and handling of a large number of system components and parameters. *Model Predictive Control* (MPC) arose as an appropriate control method that is able to fulfill these requirements by additionally complying strictly to specific constrains. This method is based on a continuous optimization of the desired control output, which considers all possible future states of the system within a predefined prediction horizon. However, the complexity of the underlying optimization algorithms causes MPC to be a compute- and memory-intensive process, which made it unfeasible for energy-constrained embedded implementations for a long time. Therefore, MPC often runs on dedicated,

energy-hungry computer systems instead of embedded platforms. The realization of embedded MPC systems and therewith a significantly more energy-efficient implementation poses great challenges on the embedded system design:

- In the last decades, high-performance computing tasks like MPC profited significantly from new hardware computing platforms that offer massive parallelism like FPGAs. Compared to sequential processing on a CPU, they provide a comparable performance at significantly lower operating frequencies and energy demands due to dedicated hardware structures. However, the system design holds major design challenges due to FPGA-specific characteristics, such as limited amount of high-bandwidth memory and hardware resources.
- This leads to design tradeoffs between memory bandwidth and size, degree of parallelism, needed hardware resources and real-time capability.
- Hardware-based implementations require a significantly higher and more challenging implementation effort than software-based implementations. Furthermore, they become more efficient the more they are dedicated for a specific application, and thus lose flexibility.

Contributions and Thesis Structure

The scientific approaches presented in this work have been investigated and evaluated with respect to system architecture, application and design methodology in order to manage the vast number of potential design decisions (Figure 1.6).

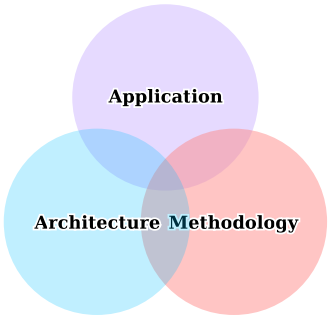

Figure 1.6: Design strategy: Investigating systems with respect to architecture, application and methodology.

All scientific investigations have been conducted in the context of appropriate embedded system applications and evaluated with dedicated implementations on real hardware. Table 1.3 shows the key aspects of these applications and in which chapters of this thesis they are discussed.

The architectures for these applications — such as computing hardware platform, memory subsystem, implementation strategy, etc. — have been worked out in particular design space explorations. Table 1.4 gives an overview of the respective design decisions.

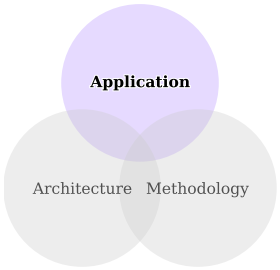

Autonomous Embedded Systems	Mobile Embedded Systems	Stationary Embedded Systems
Small-Scale, Self-Powered Sensors and Wearables *Chapter 2*	Wearable Wireless Sensor Networks *Chapter 3* Secure, TLS-based IoT Infrastructures *Chapter 4*	Embedded Control for Industrial and Infrastructure Systems *Chapter 5*

Table 1.3: The embedded system applications that have been investigated in this work and their related chapters of this thesis.

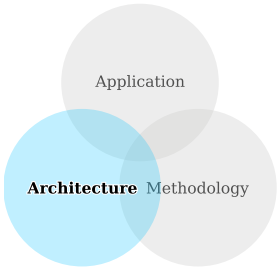

Autonomous Embedded Systems	Mobile Embedded Systems	Stationary Embedded Systems
TI MSP430FR Hardware Platform	ARM Cortex M Hardware Platform	Xilinx Zynq Hardware Platform
Kinetic & Thermal Energy Harvesting	Bluetooth Low Energy	Parallelized FPGA Implementation
Nonvolatile FRAM Memory	MEMS Sensors	

Table 1.4: Overview of the architectural properties and design decisions.

To face the challenges under the given constrains, specific design methodologies have been applied for the research. These methodologies are summarized in Table 1.5 and intensively discussed in the corresponding chapters together with the related contributions.

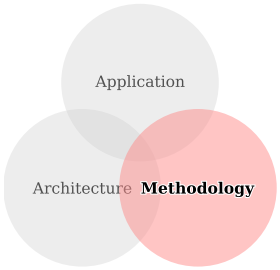

Autonomous Embedded Systems	Mobile Embedded Systems	Stationary Embedded Systems
Power-Neutral System Design	Ultra-Low Power System Design	High-Level Synthesis
Proactive Transient Computing	Perfect Forward Secrecy	Embedded Model Predictive Control
Predictive Power Loss Detection	Wireless Synchronization	Structural Data Compression
	Real-Time Processing	

Table 1.5: Overview of the applied design methodologies in the research of the addressed categories of embedded systems.

Most scientific contributions presented in this thesis have been published in conference and journal papers [31–59]. An overview of the chapters and the related key contributions is given below:

- **Chapter 2:** ***Proactive Transient Computing for Embedded Energy Harvesting Systems*** [47, 51, 58]

  A state-of-the-art approach for ensuring forward progress despite power losses is called ckeckpointing [60–62], in which the volatile system context of the processor is preserved through power losses by using nonvolatile memories. However, related work is limited to *reactive* principles that detect power losses only on short notice, which, e.g., makes the finalization of atomic operations nearly impossible.

  In this chapter, a new *proactive* and harvester-aware transient computing is proposed, which considers the internal state of kinetic energy harvesters to predict potential power losses. Furthermore, a new predictive checkpointing method is presented that monitors the exciting physical quantity of thermoelectric harvesters to predict the future trend of the energy budget. Therewith, this work makes a *proactive* power loss handling for transient computing systems possible and shows how potential power losses can be detected up to an order of magnitude earlier than related work. Moreover, particular cold-start optimization techniques are presented in order to significantly reduce the system cold-starts. This is essential for *Power-Neutral* systems that employ short output-duration harvesters, as the cold-start period length determines whether an operation is possible at all. Whereas related work either ignored the cold-start phase in their evaluation or reported cold-starts in the range of seconds [63], this work shows how the cold-start period can be minimized to as low as 75 milliseconds.

- **Chapter 3:** ***Wearable Wireless Sensor Networks for Professional Sport Performance Analysis*** [32–36, 39–41, 46]

  Long battery lifetimes for mobile microelectronic devices can only be achieved by a comprehensive combination of embedded software and hardware (hardware-software codesign). Therefore, the system design in this work follows an application-specific cross-layer approach from embedded software down to hardware. In order to cover many of the respective design challenges within one area of application, this chapter is devoted to wearable movement sensing in professional sports, in which electronic sensor devices became more and more important for in-filed sports performance analysis [64, 65]. Small-scale sensors like *micro electromechanical systems* (MEMSs) and *Inertial Measurement Units* (IMUs) have shown to be suitable to analyze the movements of athletes [66–68]. However, the system design of wearable sensor devices — ranging from embedded analysis algorithms down to the low power hardware design — represents great challenges due to the limited energy budget and processing power of mobile embedded devices [69].

In this chapter, a new hardware design approach of a wearable sensor device is presented, whose inertial measurement performance and accuracy is capable of competing with state-of-the-art stationary laboratory equipment of professional sport performance analysis systems. Furthermore, a new wireless synchronization method for BLE is proposed that allows for low power real-time movement sensing with multiple wearable sensor devices. By exploiting specific characteristics in the current profile of the BLE chip, a synchronization error of less than 10 microseconds has been achieved. Compared to BLE synchronization methods in related work [70–73], this work represents the lowest synchronization error at the time of research.

- **Chapter 4:** ***Low Power High-End Security for the Internet of Things*** [52–55, 59]

  Advanced communication security protocols like TLS rely on complex and computationally intensive cryptographic algorithms, which are known for their computational complexity and related high energy requirements [74–77]. Therefore, related work proposed lightweight security algorithms as an energy-friendly alternative for wireless bandwidth-limited and energy-constrained IoT edge devices [78–81]. However, lightweight security protocols also come with reduced security levels, which is why many state-of-the-art approaches constitute serious security risks for edge computing.

  This chapter shows how TLS can be optimized and implemented on resource-constrained, battery-powered IoT edge devices with bandwidth-limited wireless communication to ensure years of battery lifetime despite frequent communication.

- **Chapter 5:** ***Embedded Model Predictive Control for Technical Infrastructure Facilities*** [48]

  MPC is the appropriate control technique for large systems as it can handle multiple interdependencies by predicting and considering its control states in the future (prediction horizon). The evolvement of computing platforms that exhibit massive parallelism like FPGAs enabled the feasibility of MPC implementations in an embedded and low power fashion during the last decade [82–85]. However, limited amount of available hardware resources, e.g., high-bandwidth memory, on FPGAs at the time of research also limited the complexity and possible prediction horizon length of potential control tasks. Therefore, FPGA-based MPC was more or less unfeasible for environmental and smart infrastructure systems in which long prediction horizons are inevitable for a stable operation [86].

  In this chapter, a new data compression strategy is presented for an embedded MPC task on FPGAs, which allows for a prediction horizon that is significantly longer than in related work. The control task is dedicated for a smart water supply infrastructure system and the presented low power implementation enables to be embedded in the electronic circuitry of a modern water pump.

Many contributions have been worked out in collaboration with industry partners and other scientists:

- Parts of the work on *Proactive Transient Computing for Embedded Energy Harvesting Systems* have been conducted together with M.Sc. Vitor Kieling, and therefore are also discussed in his master thesis [87].
- The work on *Wearable Wireless Sensor Networks for Professional Sport Performance Analysis* has been conducted in collaboration with the Institute for Sports and Sport Science of the University of Dortmund and the department of Prof. Dr. Thomas Jaitner. Related scientific research is discussed in the dissertation of Dr. Marcus Schmidt [88] and will be discussed in the upcoming dissertation of Dipl.-Ing. Sebastian Wille.
- The work on *Low Power High-End Security for the Internet of Things* has been conducted in collaboration with KSB AG in Frankenthal and other scientists of the Microelectronic Systems Design Research Group at the University of Kaiserslautern. Further scientific research will be discussed in the upcoming dissertations of M.Sc. Frederik Lauer and M.Sc. Maximilian Schöffel.
- The work on *Embedded Model Predictive Control for Technical Infrastructure Facilities* has been funded by the German Ministry for Economic Affairs and Energy (funding no. 03ET1278C) and conducted in collaboration with the Department of Control Systems at the University of Kaiserslautern, held by Prof. Dr.-Ing. Steven Liu. The associated scientific research from the control systems perspective is discussed in the dissertation of Dr.-Ing. Felix Berkel [89].

There are further contributions that have been published, but which are not explicitly discussed in this thesis:

- Analysis of the retention behavior of *Double Data Rate* (DDR)3 and DDR4 *Dynamic Random Access Memory* (DRAM) *Dual Inline Memory Modules* (DIMMs) [38, 43–45].
- Capacitive monitoring of human nutrition [31].
- New education methods for physics lessons with augmented reality [50, 56].
- Evaluation of embedded machine learning on low power GPUs [49].
- New system concepts for a precise, *radio frequency* (RF)-based indoor localization system [37, 42].
- New methods to incorporate miniaturized sensor units into machine elements for the monitoring of the operational state and oil quality [57].

Outline

As shown in Figure 1.7, this thesis consists of seven chapters, of which chapter 2 to 5 discuss energy-related design challenges of embedded systems and present new scientific contributions. Each of these chapters focuses on a specific category of embedded systems as defined in Table 1.2. Chapter 6 presents a conclusion to the previous chapters and points out future work. And chapter 7 presents a German summary of this thesis.

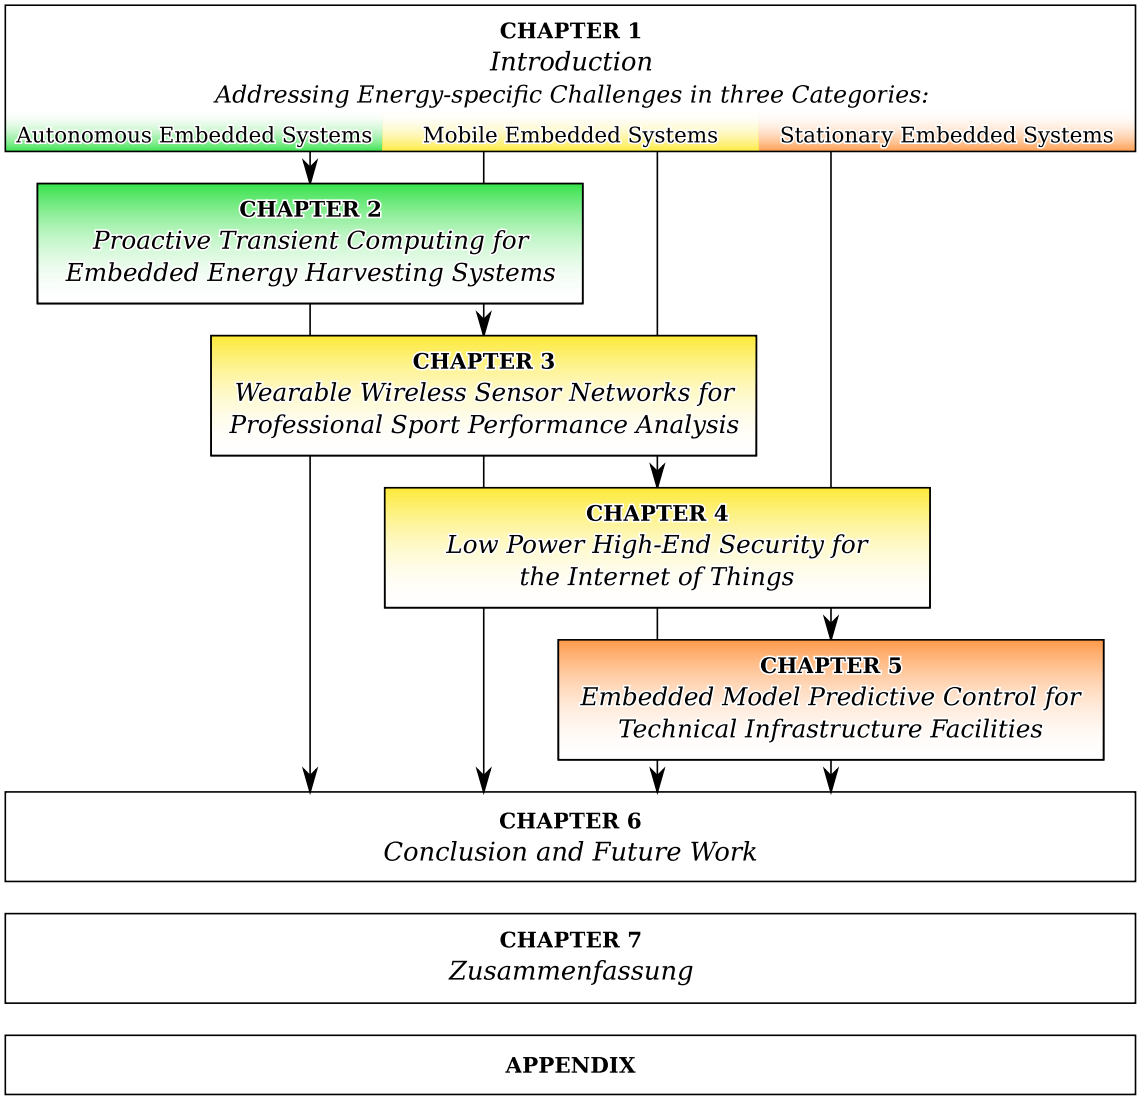

Figure 1.7: Structure and outline of this thesis.

Chapter 2

Proactive Transient Computing for Embedded Energy Harvesting Systems

”Inmitten der Schwierigkeiten liegt die Möglichkeit.”

Albert Einstein

This chapter addresses energy-specific design challenges of autonomous embedded systems, i.e., battery-free devices that scavenge their entire operational energy from the environment using energy harvesters (Table 2.1).

Autonomous Embedded Systems	Mobile Embedded Systems	Stationary Embedded Systems

Table 2.1: This chapter addresses autonomous embedded systems that are powered by energy harvesting.

Energy harvesting is a promising design concept for powering autonomous embedded devices. But, their power output is usually unstable and transient. Without proper measures, this leads to unpredictable power losses for the load system. As far as conventional embedded systems are concerned, this results in frequent restarts of the *Microcontroller Unit* (MCU) and hinders the application to make a continuous forward progress. A well-known strategy to prevent the system from being at the mercy to such power losses is to employ an energy buffer. The harvested energy is stored and the load system spends most of the time in power-down mode. Once enough energy has been collected, the system can start its operation and perform a reliable execution due to a dependable amount of energy in the buffer. This concept is known as *Energy-Neutral* system design, as the system as a whole only consumes the energy that has been harvested, no additional energy needs to be fed into the system [90–93]. However, the energy buffer involves several drawbacks, such as specific space occupation, cost and efficiency degradation [94]. The latter drawback comes

from the fact that common electrochemical energy buffers, such as rechargeable batteries, suffer from aging effects after experiencing many charge and discharge cycle. As energy harvesting-powered systems often show many charge cycles within a small amount of time, aging effects are likely to affect the system efficiency early after system deployment.

An alternative concept that completely omits energy buffers is known as *Power-Neutral* system design, as the harvested power is not buffered over time but consumed instantly [95]. As a result, these systems experience frequent and mostly unpredictable power losses. Therefore, they are called transient computing or intermittently powered systems.

These systems can be divided into state-of-the-art *reactive* and *proactive* transient computing systems. *Reactive* systems detect upcoming power losses only at very short notice, i.e., in the range of microseconds. Therefore, they must react to the upcoming power loss immediately and stop all regular execution. This leads to strictly limited application possibilities, as, e.g., a reliable finalization of atomic operations becomes nearly impossible. In contrast, *proactive* systems exploit physical interrelations of the transduction in the energy harvester as well as physical properties of its exciting physical quantity to predict possible power losses significantly earlier, i.e., in the range of milliseconds to seconds. Whereas *reactive* systems are limited to applications, where sudden interruptions of the code execution can be tolerated, *proactive* systems are promising candidates for replacing many conventional embedded systems by a more sustainable solution. This thesis focuses on *proactive* transient computing systems and faces some of the new design challenges, which have not been addressed by related work before.

This chapter is structured as follows: Section 2.1 covers the technical background of transient computing systems and highlights the differences to conventional embedded systems design. Furthermore, the specific design challenges are discussed together with related work. Section 2.2 presents the new scientific contributions for *Power-Neutral* transient computing systems.

2.1 Technical Background, Scientific Challenges and Related Work

Many design aspects and challenges for transient computing systems are different from those of conventional microelectronic systems. Figure 2.1 shows the design space that contains the relevant technological aspects of *Power-Neutral* transient computing systems.

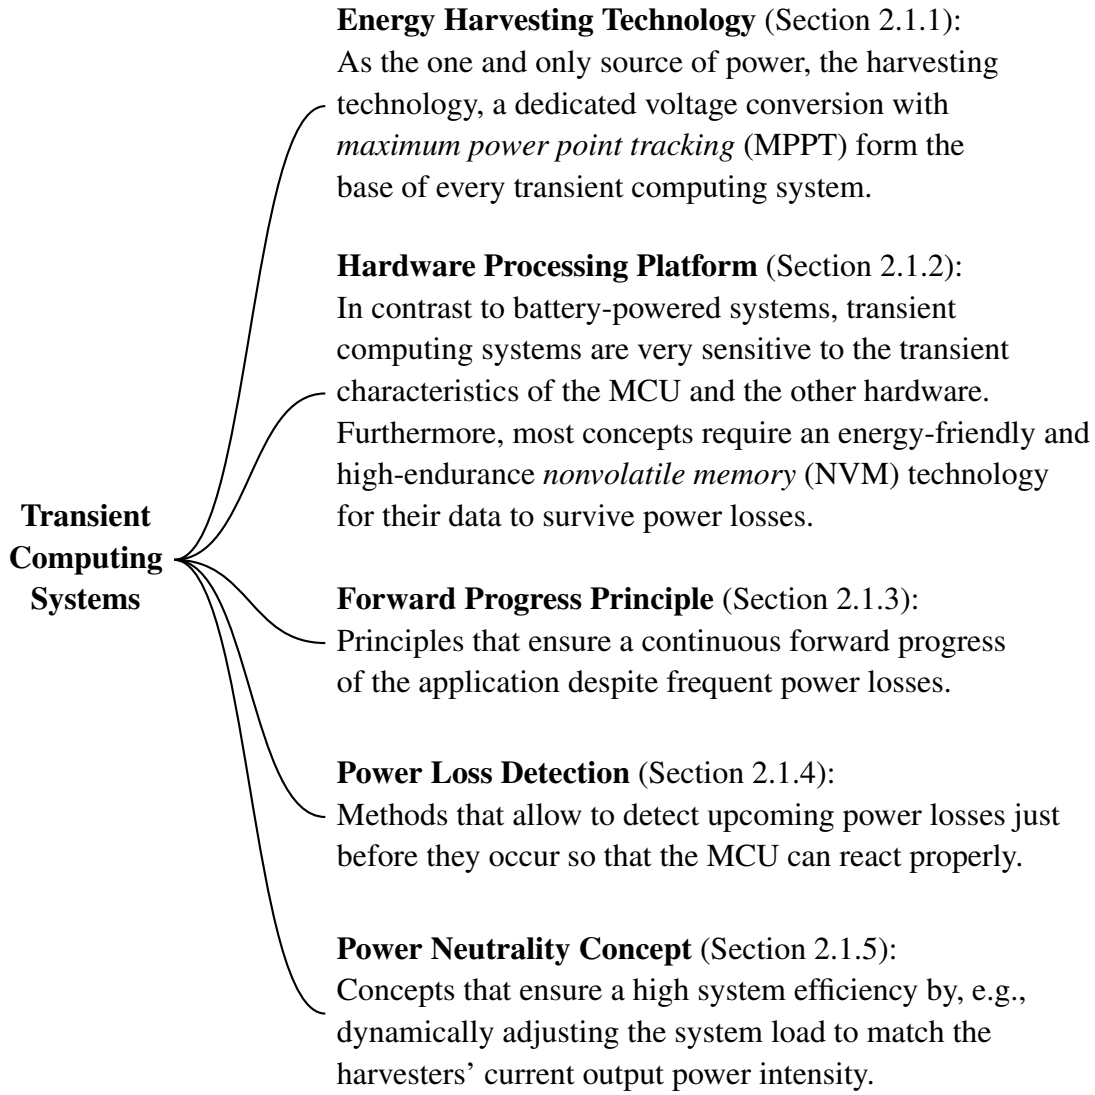

Figure 2.1: Designs space and technological aspects of *Power-Neutral* transient computing systems.

The shown design space equally applies to *reactive* and *proactive* systems. The distinction between *reactive* and *proactive* systems is mainly reflected in the applied power loss detection method and the subsequent power loss handling, as discussed in the respective section.

2.1.1 Energy Harvesting Technology

The basic concept of energy harvesting is to scavenge energy from environment and transform it to electrical energy. Therefore, energy harvesters or transducers exploit specific physical effects, e.g., the photo electric effect, to generate electricity out of the targeted energy source, e.g., light. More technical details can be found in the appendix Chapter B.2. Figure 2.2 gives an overview about different environmental sources of energy harvesting and related work.

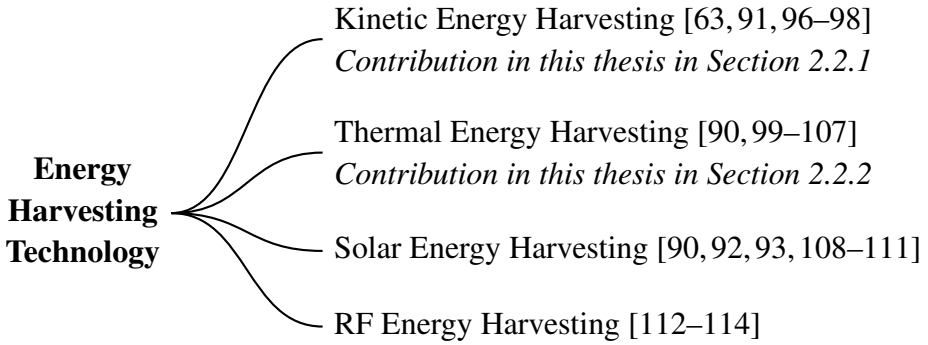

Figure 2.2: Environmental sources of energy harvesting and overview of related work.

Table 2.2 shows typical transduction principles of energy harvesting and their basic principle of function.

Source	Transduction	Signal	Principle	Requirement
Kinetic	Piezoelectric	AC	Apply compression on crystalline materials	Motions or vibrations
	Electromagnetic	AC	Change of magnetic field under movement	
	Electrostatic	AC	Change of electric field under movement	
	Triboelectric	AC	Frictional contact of two different materials	
Thermal	Thermoelectric	DC	Spatial temperature gradient	Temperature difference
	Pyroelectric	AC	Temporal temperature difference	
Solar	Photovoltaic	DC	Convert light into electricity	Bright environment
RF	RF radiation	AC	Convert electromagnetic wave into electricity	Radio coverage

Table 2.2: Typical energy harvesting sources and techniques [115].

The electrical characteristics of an energy harvester, i.e., its output voltage range and form (*direct current* (DC) or *alternating current* (AC)) as well as its power output intensity and duration, depends on the size of the transducer, the intensity of the scavenged environmental energy and how the transducer exploits the underlying physical effect. Table 2.3 shows typical electrical output characteristics for small-scale energy harvesters as reported by related work.

Source	Transduction	Reported Power Output
Kinetic	Piezoelectric	$49\mu W/cm^2$ (walking, $3km/h$) [116]
	Electromagnetic	$140mW$ (low frequency vibrations) [117]
	Electrostatic	$50\mu W$ [117]
	Triboelectric	$370\mu W/cm^2$ (wind, $15m/s$) [118]
Thermal	Thermoelectric	$14\mu W/cm^2$ (on human wrist) [101]
	Pyroelectric	$10nW$ (per chemical reaction) [119]
Solar	Photovoltaic	$7mW/cm^2$ (at $128klux$) [120] to $26.7mW/cm^2$ (at $128klux$) [121]
RF	RF radiation	$100\mu W$ (2 feet distance to Wi-Fi router) [113]

Table 2.3: Typical electrical output characteristics for small-scale energy harvesters [115].

Due to these distinct characteristics, different harvesters require different evaluation circuits to make their output usable for an embedded system. Typical components of such circuits are listed below:

- Rectifier, in case the harvester has an AC output
- Voltage converter, such as a buck/boost converter, in order to convert the varying harvester output voltage to a constant MCU-compliant supply voltage.

As shown in Figure 2.3, these circuits impose different efficiency losses along the energy flow, all the way from environmental energy scavenging up to providing a stable voltage for the load system. A major design challenge in implementing an energy harvesting system is to minimize these conversion losses and to reduce the cold-start overheads.

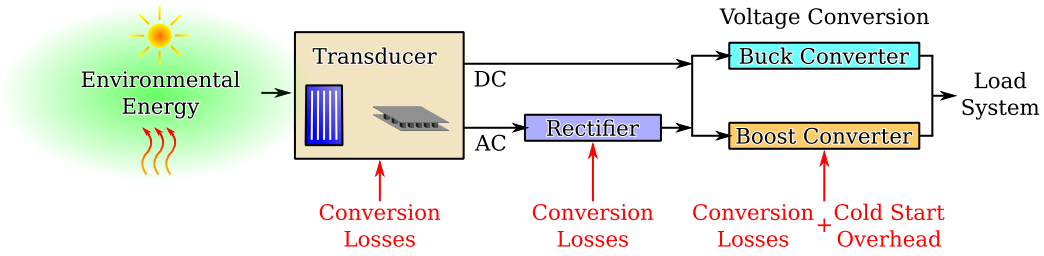

Figure 2.3: Typical losses along the energy flow in the energy harvesting subsystem.

In related work, energy harvesting has been employed as power supply technology for a long time, especially in, according to the application, obvious reasons such as solar cell-powered outdoor wildlife monitoring [93]. In the recent years, it has gained a lot of attraction for the upcoming application field of wearable electronic devices. Specifically, those harvesting technologies that scavenge energy from human motion [99, 122], e.g., kinetic [63, 91, 96] and piezoelectric harvesters [117], or from human body temperature [101–104]. With InfiniTime [90], a wearable was proposed that utilized human body heat combined with small solar cells. Besides wearables, human body heat has specifically shown to power implantable medical devices [100]. Other works employed solar

cells to harvest energy from the indoor and outdoor light [108, 110], or even proprietary harvester setups like ferro-electret insoles for a step counter [123]. Many wearable system designs targeted the emerging health care field and *Internet of Things* (IoT), in which energy harvesting was used to facilitate the idea of autonomously powered devices. Such devices continuously monitor parameters like skin temperature [92], ultra violet (UV) irradiance [108] or heart rate [109]. In contrast to classical energy harvesting approaches, the field of wearable devices holds many new design challenges, mostly due to the restrictions with respect to size, weight and strictly limited output power [99].

Most of the discussed works follow an *Energy-Neutral* design approach, since they either use a super capacitor or a small rechargeable Lithium battery to hide the transiency of the harvester source to the load system [63, 90–93, 96, 109, 110]. Thus, their relevance to this thesis is limited to the energy harvesting technology. One of the first battery-less, energy harvesting-powered applications was the Wireless Identification and Sensing Platform (WISP) [112], which is based on *radio frequency* (RF) harvesting and proposed for *radio frequency identification* (RFID) applications.

Ma et al. [115] gave a comprehensive overview of related work about energy harvesting-based developments for mobile microelectronic systems that also considers intermittent computing approaches.

2.1.2 Hardware Processing Platform

In conventional embedded systems, the hardware processing platform is usually based on an MCU that fulfills the system requirements in terms of performance characteristics and peripheral features. However, regarding *Power-Neutral* transient computing systems, there are additional specific requirements imposed on the processing platform.

Specific Requirements

Since energy harvesters typically come with low conversion efficiencies and thus low power outputs, a low power consumption is essential. Whereas most MCUs are optimized for low average energy demands, the peak power demand is the most relevant characteristic for *Power-Neutral* transient computing systems. This peak demand must not exceed the capabilities of the respective harvester, which might be at microwatt levels (Table 2.3).

High flexibilities with respect to the tradeoff between computing performance and power demand allow for a fast and dynamic scaling of the power demand. This is important in order to adjust the load to the highly varying power output of the harvester. The matching between harvester output and system load determines the degree of power neutrality, and thus the overall system efficiency (see Section 2.1.5).

For conventional embedded systems design, the transient characteristics of the processing hardware are negligible. However, in transient computing systems, frequent power

losses and subsequent power returns cause plenty of cold-starts of the entire hardware. Thus, the transient characteristics such as start-up energy, oscillator settling time and initialization delay have a significant influence on the system efficiency and the relationship between the time-limited power availability and forward progress.

In order to conserve the volatile system data through power losses, a *nonvolatile memory* (NVM) is required (see Section 2.1.3). It is important to employ high-endurance types to guarantee that the lifetime of the product is not limited and not susceptible to the total number of power losses. NVMs based on Flash or *electrically erasable programmable read-only memory* (EEPROM), e.g., restrict the number of possible power losses due to their low endurance such that their applicability for transient computing systems is limited. In contrast, emerging technologies like *ferro-electric RAM* (FRAM) or *Magnetoresistive RAM* (MRAM) come with almost unlimited endurance. More information about different state-of-the-art NVM technologies can be found in the appendix Chapter B.3.

Appropriate Platforms

Due to the specific requirements for the hardware processing platform, different hardware architectures have been applied for transient computing systems by related work as shown in Figure 2.4.

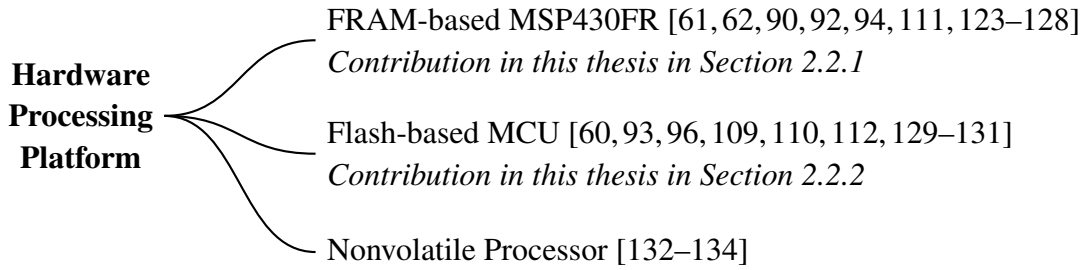

Figure 2.4: Different hardware processing platforms for transient computing systems and overview of related work.

Early works mostly employed conventional MCUs with embedded flash memory [112]. As these approaches used this flash memory for *Checkpointing* (CP) [60] (Section 2.1.3), they experienced significant backup overheads compared to recent works that employed more advanced NVM technologies.

With the MSP430FR series [135], Texas Instruments has developed a commercially available 16-bit MCU that incorporates nonvolatile FRAM on chip. This FRAM can be configured as unified memory, which means that it can be used as program memory and regular *random access memory* (RAM), omitting the requirement for data backup of the RAM content. Due to this high degree of on board nonvolatility, the MSP430FR series has broadly been applied to transient computing systems in research for years. Besides its nearly unlimited endurance, FRAM has shown to offer up to 1000 times faster write

capability with at the same time up to 100 times lower power demand compared to Flash-based MCU systems [136]. To support transient computing on their devices, TI offers a compute through power loss (CTPL) library that provides an *application programming interface* (API) to dynamically backup and restore the volatile state of specific peripherals on chip. The popular works Hibernus [62] and QuickRecall [61] proposed two pioneering CP approaches for the MSP430FR series, which mainly differ in the memory mapping. Hibernus uses the conventional *static random access memory* (SRAM) memory as main working memory and backups all the data to FRAM just before a power loss. QuickRecall in contrast, utilized the ability to use the FRAM as unified memory, i.e., substituting the SRAM with a specific address range of the FRAM. This way, the CP backup process does not have to back up the SRAM content. However, since FRAM has a lower energy efficiency than SRAM, this approach results in a higher energy demand during operation than Hibernus [124]. On the other hand, QuickRecall outperforms Hibernus for high frequencies of power interruption [124]. However, a general rule cannot easily be derived. Kim et al. [125] conducted a comprehensive analysis about how the memory mapping affects the energy consumption of FRAM-based embedded devices. They conclude that an average of 50% of energy savings can be achieved, when the memory mapping is specifically adjusted to the specific application.

Besides FRAM-based MCUs, the popular flash-based MCU ARM Cortex M series have been subject of several studies. Balsamo et al. [129, 130] evaluated the integration of Hibernus into ARM *mbed* [137], a popular programming framework for ARM controllers. The implementation required an additional capacitance of $1000\mu F$ [129] up to almost $5000\mu F$ [130] to ensure a reliable CP backup procedure, whereas for the FRAM-based implementation only $16\mu F$ was sufficient [62]. Besides the significantly higher energy demand for programming flash memory, one reason for this is that flash requires an energy-hungry erase procedure prior to programming.

A rather unconventional type of NVM has been used by Ghodsi et al. [138] who evaluated their optimal CP approach on a system that employed *Resistive RAM* (RRAM).

The use of *Nonvolatile Processors* (NVPs) [133] in transient computing systems is a hardware-based solution that theoretically omits the need of CP as all memory cells inside this type of processor are realized using NVM technologies [132]. However, compared to SRAM, which is the state-of-the-art technology for processor-internal memory cells, most NVM technologies impose energy overheads [124, 126, 139] and special hardware requirements [140]. At the time of this research, NVPs are not yet commercially available [94] [126] [134], and remain a theoretical research topic with limited relevance for practical transient computing systems.

2.1.3 Forward Progress Principle

Central Processing Units (CPUs) and MCUs require a minimum voltage $V_{CPU\ min}$ to operate. In the event of a power loss, i.e., once the supply voltage V_{CC} drops below this minimum, the execution of instructions stops immediately and all volatile data is lost or at least corrupted. On power return, usually a power-on reset is issued and the MCU starts over from scratch by executing the first instruction in the program code. Thus, frequent power losses cause a perpetual restarting of the MCU and hamper the application from making a steady forward progress. Therefore, technical concepts that guarantee a forward progress for transiently powered devices despite frequent power losses are essential. Figure 2.5 shows the state-of-the-art forward progress principles that have been employed by related work.

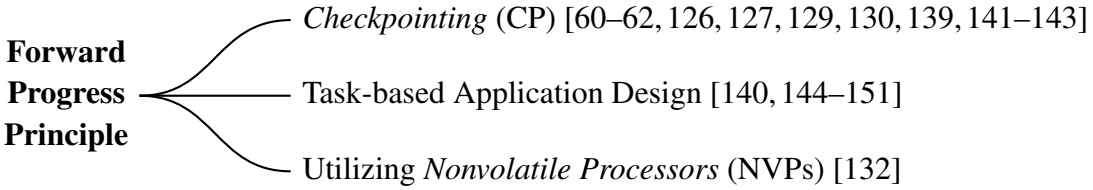

Figure 2.5: Different forward progress principles for transient computing systems and overview of related work.

Checkpointing

The basic idea of *Checkpointing* (CP) is to back up the volatile context, e.g., registers, stack and RAM content, of the MCU to an NVM just before a power loss. After power return, the context is restored from the NVM and the MCU continues its code execution where it stopped before. Consequently, CP requires some sort of NVM and in most cases a *power loss detection* (PLD) (Section 2.1.4) method that detects power losses before they actually occur. CP can be done on demand, i.e., just before a power loss, or at predefined checkpoints in the application. On-demand CP — also known as *dynamic* CP — requires a continuously working PLD method in order to detect upcoming power losses, e.g., by detecting when V_{CC} deceeds a threshold V_{th}. This procedure is shown in Figure 2.6, where the CPU activity is illustrated to be dependent on the voltage supply level V_{MCU} and split into restore and backup procedures before and after normal execution. The threshold V_{th} is split into the voltage levels V_{th-} and V_{th+} to ensure a hysteresis between PLD and identification of a power return. The continuous PLD results in a continuous overhead, but in return only issues CP backup and restore operations when required, thus saving overheads due to futile NVM write operations. Useless operations are only issued in the seldom case, where V_{MCU} does not fall below V_{MCUmin} after falling below V_{th-} (marked as *) in Figure 2.6). In contrast, CP with predefined checkpoints — also known as *static* CP — saves

the MCU context frequently, e.g., at each loop iteration or function call. Therefore, all results that have been computed after the last checkpoint are lost once a power loss occurs. This method requires a careful design to be resilient against corrupted checkpoints, as the power loss can occur during the backup process. It also causes an overhead due to many useless NVM write operations, but, in return, does not require a continuous overhead for PLD. An advanced approach is to activate a specific PLD method at the predefined checkpoints, and only issue CP when the power is likely to be lost until the next checkpoint. The energy overhead of CP depends on the employed PLD method and the energy demand of the NVM read and write operations. The latter strongly depends on the type of NVM and the amount of data that must be preserved.

Various CP methods were proposed for transient computing systems to preserve the volatile CPU state through power loss using NVMs [60–62]. Besides the CPU state, some works specifically addressed the challenges of preserving the context of peripherals [127, 141, 142].[1] In Mementos [60] the checkpoints are placed inside the application code at specific positions. A mathematical optimization of the checkpoint positioning in the application code has been presented with FlexiCheck [143]. In contrast, QuickRecall [61] and Hibernus [62] proposed to perform CP on demand, i.e., at run-time just before a power loss is about to occur. On demand CP enables a sufficient exploitation of the available energy as it maximizes the execution time and thus forward progress. But it leads to interruptions of the application at any position, so that proper termination of specific tasks cannot be ensured. As a solution, related work proposed to accumulate energy in advance for proper task termination [111, 145, 152]. But these proposals result in an *Energy-Neutral* system concept, which, however, has been valued as inevitable energy-intensive tasks [150].

Task-based Application Design

A task-based application design can help to significantly reduce the amount of data that must be saved to NVM before power loss. Whereas on-demand CP usually requires to back up the entire volatile context of the MCU, task-based design aims at only preserving task results and program pointers. A simple but representative example is an application that uses no global variables and a lot of function calls. Here, upon each new function call, only a little RAM data is required from the previous function. However, the advantage of fewer amount of backup data comes at the cost that tasks, who have been interrupted must always be restarted from scratch as shown in Figure 2.6. Furthermore, task-based approaches tend to needless backups (marked as *) in Figure 2.6), due to more frequently appearing backup procedures compared to CP. Therefore, the task granularity must carefully be matched to the expected power availability periods.

Early proposals like Dino [147], Ratchet [151] and Chain [148] relied on manual CP instruction insertions in the code, which comes with the drawback of obliviousness to the

[1] This paragraph has already been published in [51]

available energy [143]. A more advanced principle is to consider the transiency of the power supply like the task decomposition proposed in CleanCut [144]. Dedicated task control systems that execute tasks when the energy conditions are matching, can help to achieve a higher energy utilization for storage-less sensor nodes [146]. A specific method is to annotate the tasks with their respective energy and run-time requirements [145]. Other approaches are anytime algorithms concepts for transiently powered devices. Instead of aiming at finishing tasks, the respective computational results are successively approached depending on the available energy. Thus, at power losses the current result is taken and computed towards higher accuracy in the consecutive power cycle [153].

As task-based CP operates more or less agnostic to upcoming power losses, inconsistencies and in the NVM, e.g., due to interrupted CP backup write procedures, are a common problem. Proposed solutions range from multiversioning of inter task data [148] over ensuring idempotence of the task stream on hardware [149] to privatization of the task-shared data [140].

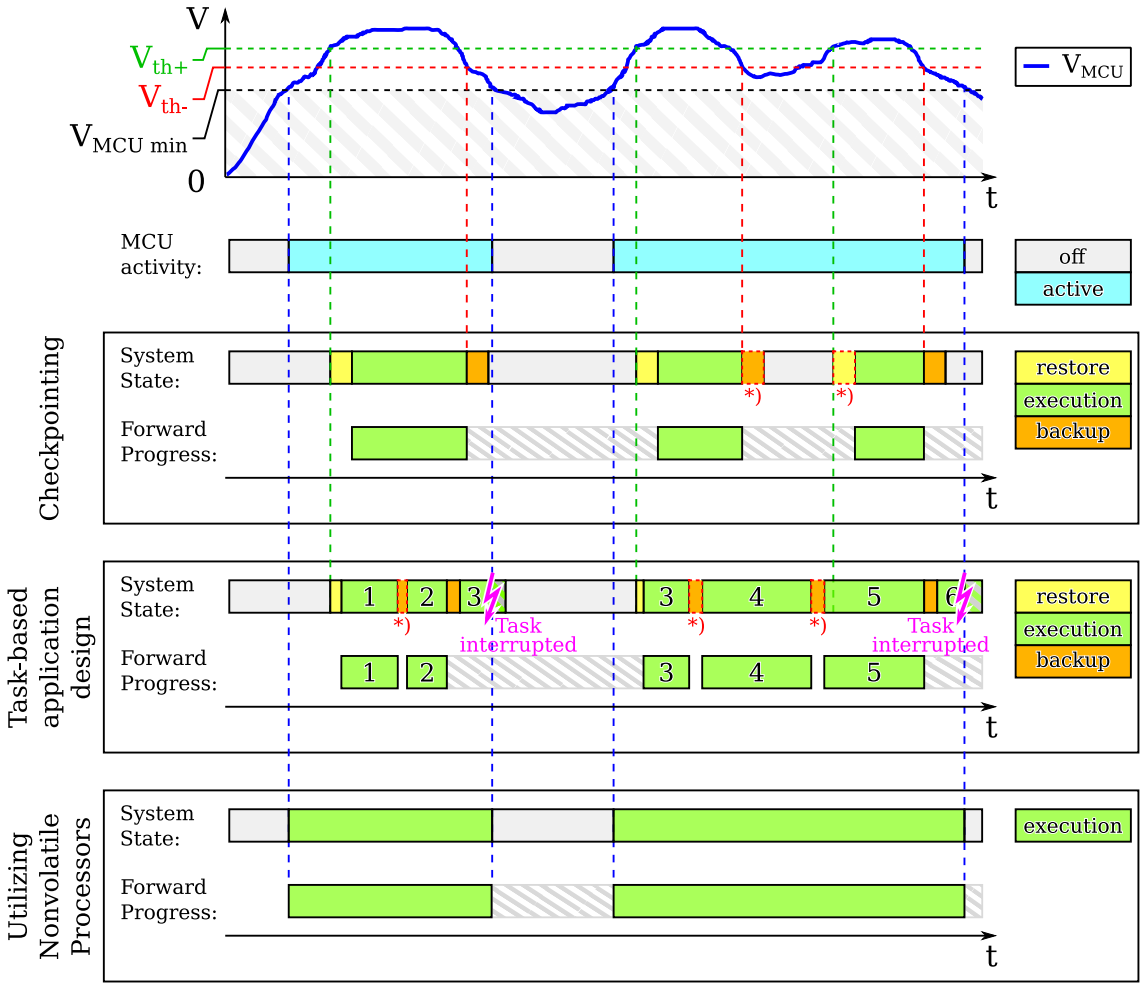

Figure 2.6: Comparison of the common forward progress principles.

Utilizing Nonvolatile Processors

A *Nonvolatile Processor* (NVP) is a hardware-based concept that incorporates only nonvolatile memory cells inside the MCU [133]. Therefore, NVPs omit the necessity of the forward progress principle methods discussed above, as nonvolatile memory cell contents do not need to be preserved through power losses. However, as discussed in Section 2.1.2, they are not commercially available at the time of this research and thus are not further considered in this thesis. For comparability purposes, Figure 2.6 incorporates the theoretical forward progress principle of NVPs during power losses.

2.1.4 Power Loss Detection

A central challenge of transiently powered devices that employ CP is to detect upcoming power losses just before they occur. At the point in time of the detection, a minimal residual energy must be left to allow the MCU to back up its volatile context to the NVM. Figure 2.7 gives an overview about different *power loss detection* (PLD) methods, which can be divided into *reactive* and *proactive* principles as discussed in the following.

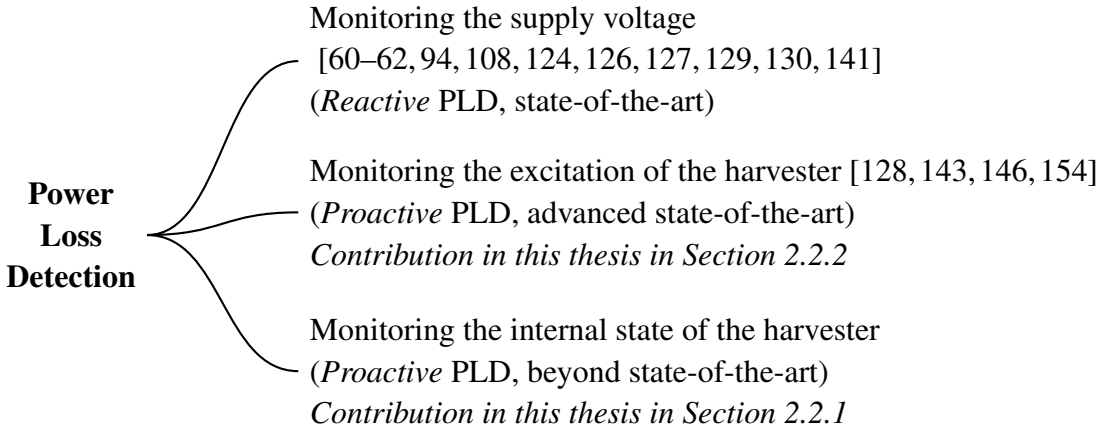

Figure 2.7: Different power loss detection methods for transient computing systems and overview of related work.

Monitoring the Supply Voltage

The state-of-the-art approach is to continuously monitor the supply voltage and to identify any voltage dropping that might indicate an upcoming power loss. Due to the decoupling capacitance, there is a specific energy left in the system once the supply voltage drops. Until it reaches the minimum operating voltage of the MCU, this energy can be used for CP backup. Depending on the respective energy demand, additional capacity must be added to the system to ensure sufficient backup energy. Figure 2.8 shows how the voltage-based

power loss detection functions during a power loss, where MCU represents the load system, C_D the overall decoupling capacitance and G the energy harvester that is protected by a diode from drawing current from the system when it is inactive.

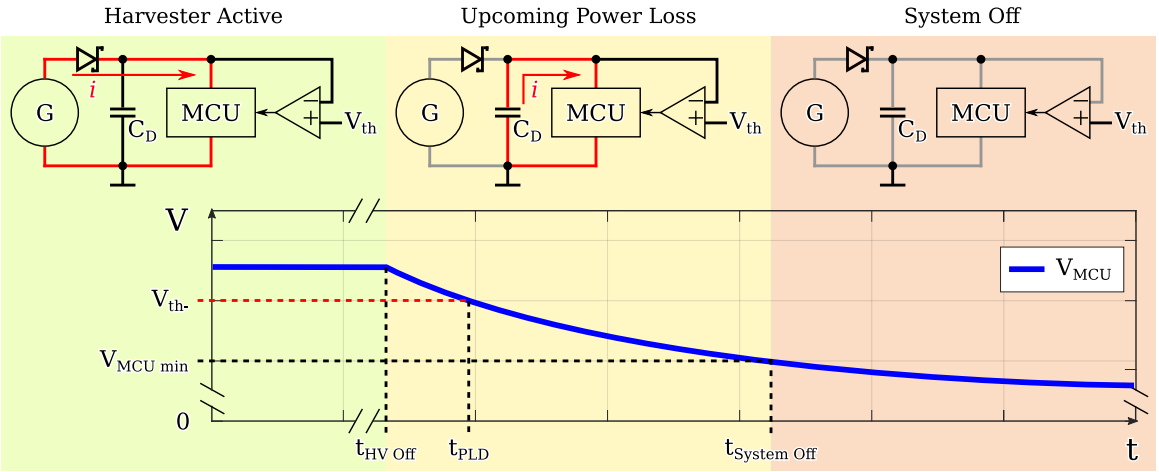

Figure 2.8: Power Loss Detection by monitoring the supply voltage.[2]

At the time where the harvester is active the load current i is provided solely by the harvester. In the event of an upcoming power loss, i.e., when the harvester became inactive at $t_{HV\ Off}$, the load current i is provided by the decoupling capacitance C_D and the supply voltage V_{MCU} drops as C_D is discharged. V_{MCU} is monitored by a comparator that recognizes when voltage drops below the threshold voltage V_{th-} at t_{PLD}. This is the time where the MCU detects the power loss and must issue the backup procedure and finalize it before V_{MCU} drops below the minimum operating voltage $V_{MCU\ min}$ at $t_{System\ Off}$.

Considering the discharge behavior of a capacitor $v_C(t) = V_0 * e^{-\frac{t}{\tau}}$, the remaining power supply duration time $t_{PSD} = t_{System\ off} - t_{PLD}$ can be calculated by

$$t_{PSD} = -\frac{V_{CC}}{i} * C_D * ln\left(\frac{V_{MCU\ min}}{V_{th}}\right). \tag{2.1}$$

To enable a fast system start-up, the decoupling capacity is usually kept as small as possible for *Power-Neutral* transient computing systems. According to Equation 2.1, t_{PSD} is usually small, which will be illustrated by the following example calculation: The MSP430FR5869 MCU [155] by Texas Instruments is a typical representative of the widely employed FRAM-based MSP430FR series. Texas Instruments recommends a minimum decoupling capacity of $C_D = 2.2\mu F$. By using characteristic values reported in the data sheet, i.e., an active operating current of about $i = 1mA$, a supply voltage of $V_{CC} = 3.0V$ and a minimum operating voltage of $V_{MCU\ min} = 1.85V$, together with the PLD threshold value used in Hibernus [62], $V_{th} = 2.17V$, t_{PSD} is determined to:

[2] Graphic has already been published in [47,51]

$$t_{PSD} = -\frac{3.0V}{1mA} * 2.2\mu F * ln\left(\frac{1.85V}{2.17V}\right) \approx 1ms \tag{2.2}$$

Equation 2.2 indicates that state-of-the-art voltage-based PLD methods require the MCU to react immediately and initiate the CP backup procedure, which must be finished within 1 millisecond in the given example. This is why these methods are referred to as *reactive* PLD methods in this work.

In Mementos [60], PLD was performed at predefined checkpoint positions, at which an analog to digital converter determines, whether the supply voltage is below a given threshold or not. Hibernus [62] and QuickRecall [61] proposed to perform CP on demand. To detect upcoming power losses, the supply voltage is monitored continuously by comparators with static thresholds. This way however, the PLD cannot be adapted to the dynamics of the harvester [124]. As a solution, Hibernus++ [94] proposed to adapt the PLD at runtime by introducing a self-calibrating circuit that adjusts the threshold to the properties of the harvester and the load. However, all discussed PLD methods, including those in other CP proposals [139, 141], basically identify power losses when the supply voltage falls below a given threshold. Such voltage-based methods only provide a PLD in the short run, thus only enabling a *reactive* transient computing. Where applicable, these methods require additional capacitances for CP that cause the system start-up to decelerate.[3]

Reactive transient computing may be sufficient for systems that employ advanced NVM technologies like FRAM that allow for a backup procedure in a very short time and with very little energy demand [124, 136]. However, especially regarding most IoT sensor devices, most systems are based on MCUs and *System on Chips* (SoCs) that only employ conventional Flash memory as NVM. Partly because Flash memory requires a block erase procedure prior to a writing operation, it requires far more energy and time for a CP backup than, e.g., FRAM [156]. Therefore, voltage-based methods always require large additional chemical energy buffers for a successful CP backup procedure [129, 157].

Monitoring the Excitation of the Harvester

An advanced state-of-the-art approach is to monitor the excitation of the harvester, i.e., its exciting physical quantity, to predict the harvested power levels. This is an emerging concept towards a *proactive* transient computing and a promising *Power-Neutral* alternative to additional energy buffers. As an example, related work investigated the prediction of energy for solar harvesters [128]. They used voltage thresholds to decide whether an intensive task can be finished or not. In [146], past solar radiation data was considered, and in [154], variations in the historic data were used to predict short-term variations. The authors of [143] utilized specific knowledge of vibration patterns of machines to predict the energy of a vibrational energy harvester.[4]

[3] This paragraph has already been published in [47]

[4] This paragraph has already been published in [58]

In this thesis, a new method for thermoelectric energy harvesters is presented in Section 2.2.2, in which the temperature gradient across a *thermoelectric generator* (TEG) is monitored by two temperature sensors in order to predict the expectable output power level.

Monitoring the Internal State of the Harvester

A new *proactive* approach, which has not been explored by related work yet, is to monitor the internal state of the harvester in order to predict upcoming power losses earlier than by monitoring its output voltage.

In this thesis, this approach is presented for kinetic energy harvesters in Section 2.2.1. The internal state of the harvester — as it is reflected in the frequency of its output voltage — is exploited to determine how much longer the power supply will be available and therewith detecting upcoming power losses early on.

2.1.5 Power Neutrality Concept

In contrast to *Energy-Neutral* systems, *Power-Neutral* systems are designed to consume the power that is provided by the harvester instantly. To ensure a good overall system efficiency, i.e., good level of power neutrality, the load must be identical to the harvester's maximum output power, i.e., its *maximum power point* (MPP). This MPP is subject to continuous alteration, as the excitation of the harvester strongly depends on changes in its environment and the transiency of its respective energy source. As a solution, the load system must employ a *maximum power point tracking* (MPPT) technique that continuously measures the harvester's MPP as well as a load alteration technique that adjust its load to this MPP. However, such techniques complicate the system design and therefore most related work did not specifically address this issue but rather just kept the load system requirements below the average power output level.

Figure 2.9 gives an overview about different techniques and concepts, which have been employed by related work to achieve a good power neutrality. In transient computing systems, the quantitative success of the respective concept is referred to as level or degree of power neutrality.

The state-of-the-art implementation approach to evaluate the harvester's MPP at runtime is to monitor the supply voltage while altering the system load. This way, some sort of a closed-loop control system is realized, leading to the optimal operating point of the load, i.e., leading to a satisfying level of power neutrality. The most basic concept to adjust the system load is to employ power gating of various system peripherals. Since this shows a relatively coarse granularity in load alteration, this technique is mostly used to deal with macro-scale variations in the harvester's MPP. A finer granularity can be achieved with frequency-scaling of the MCU, which exploits the relationship between clock frequency and current demand of MCUs.

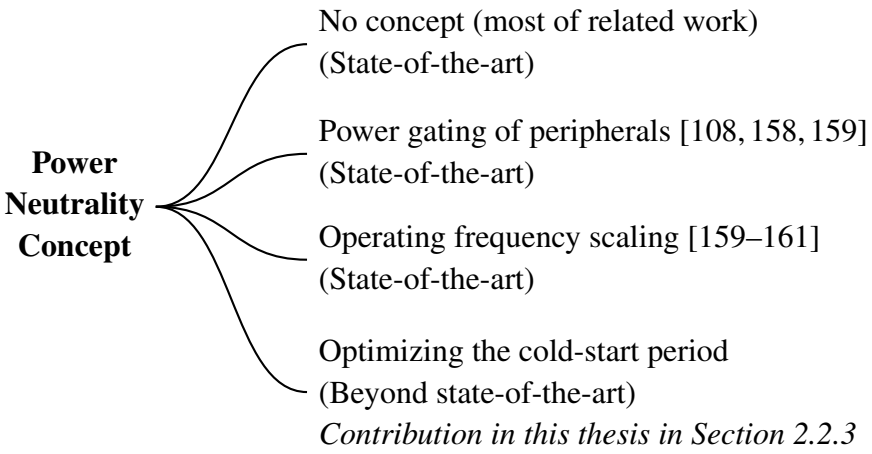

Figure 2.9: Different power neutrality concepts for transient computing systems and overview of related work.

Figure 2.10 illustrates how the discussed power neutrality concepts are used to alter the system load at run-time and shows how they impact the level of power neutrality. It shows the system efficiency as relationship between the energy harvester power output and the power profile of the load system. The level of power neutrality is defined as the relation of used energy E_{used} and wasted energy E_{wasted} to the theoretically available energy by the harvester E_{max}.

Power Gating of Peripherals and Operating Frequency Scaling

The top region of the graph illustrates the system efficiency for a constant load, i.e., where no MPPT or comparable power neutrality concept is employed. Since the load profile P_{load} is constant, the amount of wasted energy E_{wasted}, i.e., the energy that is not utilized by the system, is relatively high compared to the available power of the harvester P_{MPP}. This leads to a low level of power neutrality. The middle region of the graph shows — exemplary for three peripherals with different load characteristics — how the power gating of peripherals can improve the level of power neutrality as P_{load} can be altered during run-time to at least roughly approach P_{MPP}. However, the granularity of load alteration is limited to the usually static load characteristics of most peripherals. A more precise load alteration is shown in the bottom region of the graph, in which a dynamic scaling of the operating frequency of the system MCU (f_{MCU}) enables to change the load profile P_{load} finer-grained compared to power gating. The granularity of the load alteration depends on how precisely f_{MCU} can be set in the given MCU.

In related work, several proposals have been published to adjust the system load as accurate to the MPP of the harvester as possible. Due to the lack of an energy buffer, the implementation of MPPT like in conventional embedded systems is almost impossible for *Power-Neutral* systems [158]. Instead, the MCU must exploit its possibilities in altering its power

demand at run-time like dynamic adjustment of the operating frequency [160, 161]. Others implemented power-gating of peripherals [108, 158] or combined both techniques [159].

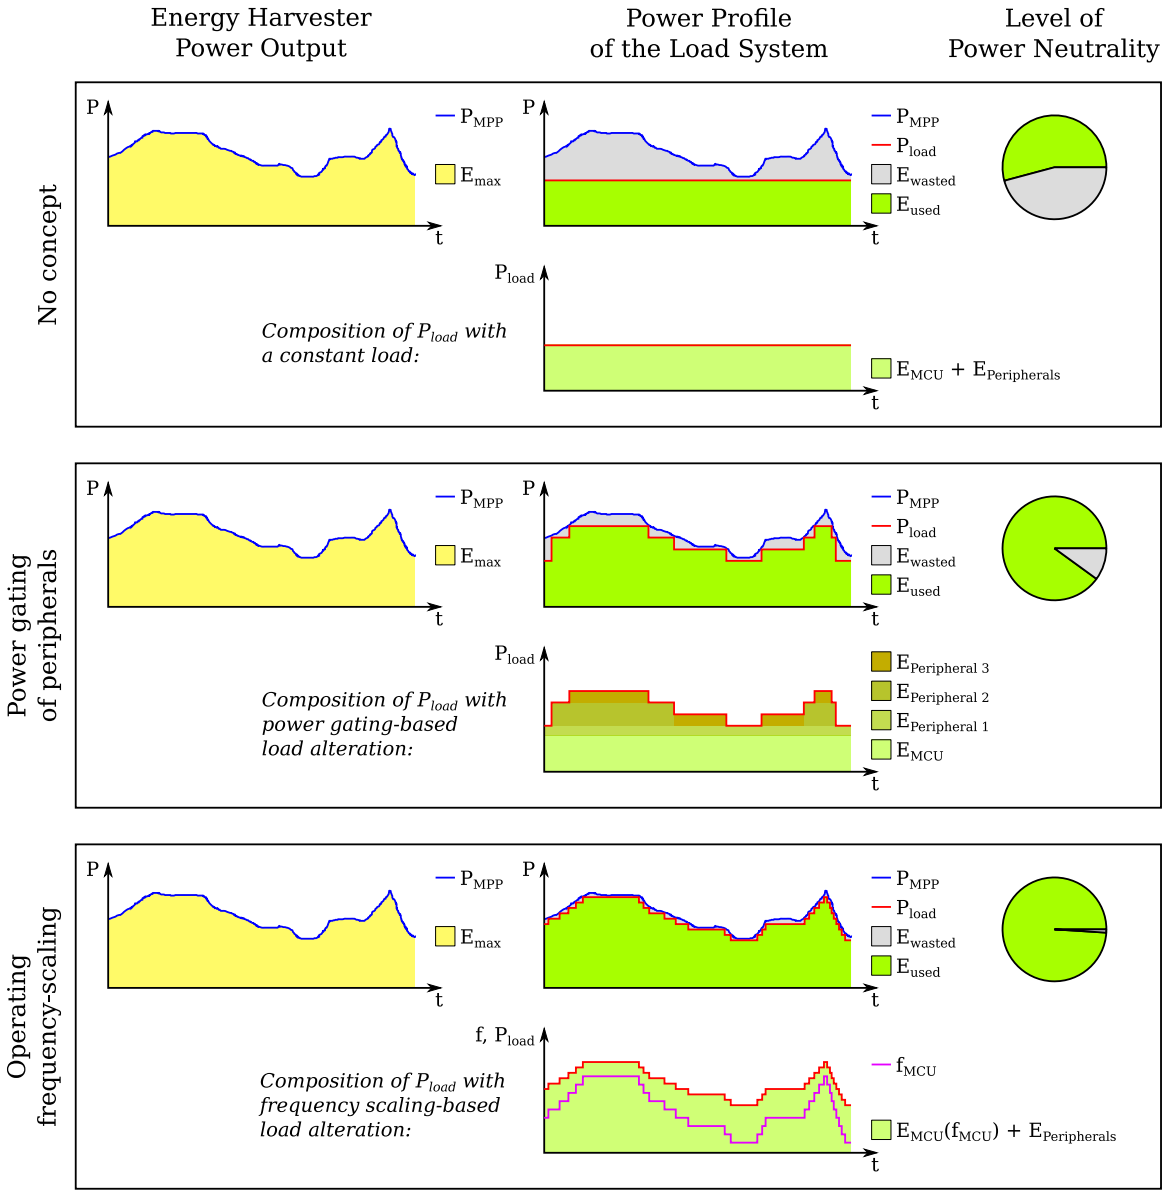

Figure 2.10: Comparison of different power neutrality concepts.

Further works did not explicitly target power neutrality, but proposed methods that help to come close to the intended objective. For instance, to optimize the overall energy efficiency of the system by a holistic analysis that considers the harvester, MPPT technique and the MCU together [162]. Or by switching the system to retention state once the harvested power becomes too weak for normal operation, which can help to prevent power failures [163].

Optimizing the Cold-Start Period

The cold-start period is defined as the time when the energy harvester starts providing electrical power to the system until the MCU is executing its first instruction. In contrast to conventional embedded systems that come with a reliable energy basis, *Power-Neutral* transient computing systems experience many power losses that result in a complete energy loss in the system and thus undergo a complete cold-start period upon each power return. Therefore, the cold-start period contributes significantly to their overall power neutrality. Despite this fact, related work has not yet specifically addressed this issue and commonly accepted it as unavoidable.

In this thesis, in Section 2.2.3, a deep analysis of the cold-start period is presented and specific optimization strategies are proposed that show how this period can be reduced significantly to improve the overall level of power neutrality.

2.2 New Contributions

In this section, the new scientific contributions of this work in the research field of *Power-Neutral* transient computing systems are presented. These contributions are divided into three topics. The first and the second address the design challenges of *proactive* systems for detecting power losses earlier than state-of-the-art *reactive* systems, and the third addresses the cold-start challenges of either system types:

- Section 2.2.1 presents how the monitoring of the internal state of a kinetic energy harvester enables a harvester-aware transient computing that allows to predict potential power losses earlier than related work. It particularly shows how this can be utilized to ensure the finalization of atomic operations, and therewith supplying a solution to a common limitation of state-of-the-art *reactive* systems.
- Section 2.2.2 presents how the monitoring of the exciting physical quantity of a thermoelectric energy harvester allows to predict its power output levels, which facilitates a predictive CP mechanism. It particularly shows how this can be utilized to schedule CP backup procedures on energy-hungry NVMs like Flash memory, and therewith presents a *Power-Neutral* alternative to large additional energy buffers as commonly required by state-of-the-art approaches.
- Section 2.2.3 presents new optimization strategies to significantly reduce the cold-start period of *reactive* and *proactive* transient computing systems and therewith improve the overall level of power neutrality compared to system designs presented by related work.

All scientific investigations have been conducted on the dedicated hardware evaluation platform *EternAlyzer*, which is described in the appendix Chapter B.1.

2.2.1 Harvester-aware Transient Computing for Kinetic Harvesters[5]

The timely detection of upcoming power losses is essential for a reliable checkpointing process. Moreover, various applications require early detections to, e.g., ensure the finalization of atomic operations. However, common voltage threshold-based methods only allow short-term detections.

In this section, a new methodology is proposed that allows early detections by exploiting physical characteristics of the harvester. To this end, a small-scale kinetic energy harvester has been considered, the MGS26.4 [164] by Kinetron (Figure 2.11), which has been used in related works for powering autonomous devices [63, 91, 96, 98, 99]. Moving the harvester stimulates its inner rotating mass (flywheel) to rotate, which in turn drives a micro electromagnetic generator that converts this rotation to electrical energy. Due to the inertia of this mass, the power output does not stop abruptly, but gradually decays after the excitation of the harvester is over.

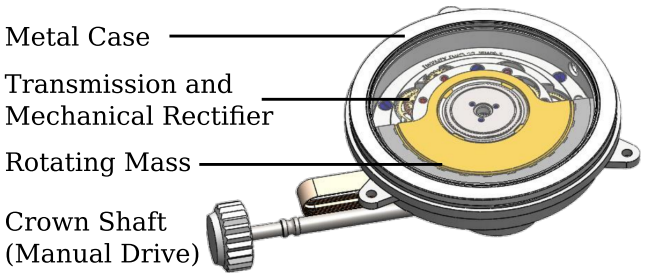

Figure 2.11: Micro Kinetic Harvester MGS26.4 [164].

Kinetic Harvester Characteristics

Figure 2.12 shows the output voltage that is generated when the harvester is placed vertically and then slowly tilted to cause the flywheel to start moving downwards. Due to its mass and the internal spring barrel, the flywheel oscillates back and forth a couple times, before it remains in its final rest position. This is reflected in the subsequent voltage bursts. Sudden and strong movements cause the output voltage and frequency to be higher and the burst duration to be longer.

The rotating speed ω increases during excitation, e.g., caused by the gravitational acceleration g as shown in Figure 2.12, and decreases afterwards. Thus, the *inflection point* (IP) of ω marks the point where the excitation is finished. After the inflection point, the harvester continues to provide power until the rotation of the flywheel decays, indicated as striped areas in Figure 2.12. Hence, the power output lasts longer than the excitation. The physical reason for this phenomenon is the kinetic energy E_{kin} that is mechanically stored in accelerated masses, which depends on the intensity of the initial acceleration. In case of the harvester, this intensity is reflected in the rotating speed ω at the inflection point.

[5]This section and the related graphics have already partially been published in [47, 51]

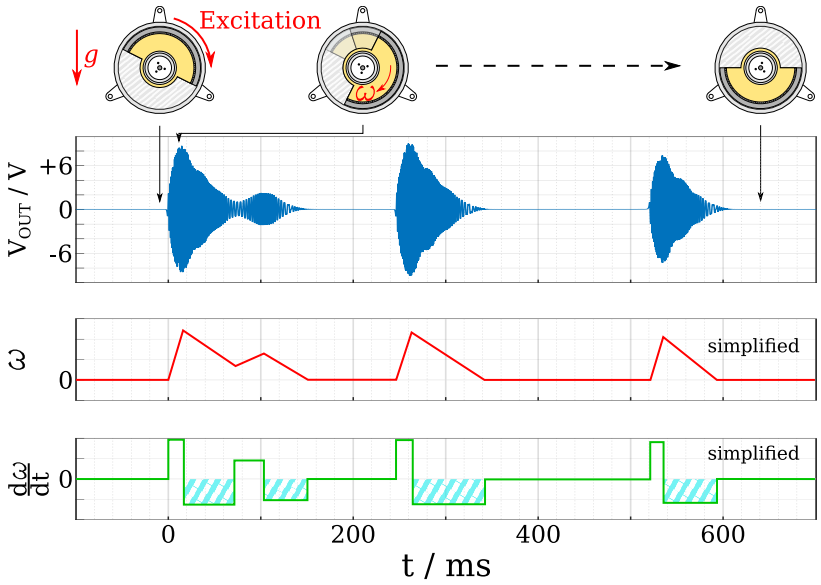

Figure 2.12: Open-loop output voltage and flywheel rotation profile of the kinetic harvester after being excited by gravitational acceleration. The rotational speed has not been measured directly, therefore the two lower graphs are only a qualitative and simplified representation.

Regarding the physical construction of the harvester, the kinetic energy can be described as rotational energy

$$E_{kin} = E_{rot}(\omega) = \frac{1}{2} * \Theta * \omega^2, \tag{2.3}$$

where Θ is the inertia torque of the flywheel with respect to the rotational axis, which depends on the mass distribution of the flywheel. The electrical energy is obtained by $E_{el} = E_{rot} - E_{loss}$, where E_{loss} contains the losses of the generator due to friction and copper resistance of the coil.

Our new methodology exploits the mechanically stored energy in the flywheel of the harvester as a predictive energy storage. In order to quantify the amount of this energy at run-time, the intensity of the initial excitation has to be determined by the application. Therefore, the rotational speed ω at the inflection point must be measured. Based on the physical construction of the harvester and the underlying physical laws, the rotational speed ω correlates with the voltage and frequency of the generator output. The transient impact of the internal spring barrel is neglected, as it does not effectively add energy. According to the law of electromagnetic induction, the output voltage $V_G = \frac{d\Phi}{dt}$ is induced by the magnetic flux Φ in the coil. Given the temporal behavior of Φ to be sine function $\Phi(t) = \Phi * sin(k_\omega \omega t)$, the amplitude of the voltage can be described as

$$V_G(\omega t) = \Phi * k_\omega * \omega * cos\left(k_\omega \omega t\right) = V_G * cos\left(k_\omega \omega t\right), \tag{2.4}$$

where k_ω represents the correlation between the rotational speed ω and the frequency of Φ, depending on the construction of the internal transmission, rotor and stator.

Equation 2.4 points out, that the rotating speed of the flywheel is reflected in the peak output voltage and the frequency equally. Thus, the rotating speed of the flywheel can be obtained by measuring the voltage or the frequency respectively. Even though the internal spring barrel mechanically decouples the rotating mass from the generator [98], our results showed a direct correlation between excitation intensity and output frequency.

System Implementation

As shown in Figure 2.13, a bypass path has been added for the frequency measurement.

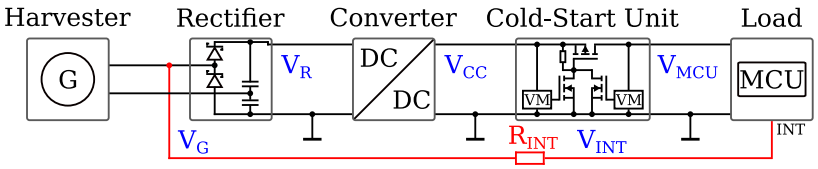

Figure 2.13: Proposed hardware structure with bypass for frequency measurement. For further details about the shown hardware components see Chapter B.1.

After R_{INT}, the output voltage of the harvester is clamped by the protection diodes of the MCU, which forms a rectangular signal at the input pin (V_{INT}). The pin is configured to fire an *interrupt request* (IRQ) at the falling edge. Inside the corresponding IRQ_f routine, the actual count of a timer is captured and compared to the previous count value. This way the frequency of V_{INT} is calculated for each period. Once the frequency values decrease, the MCU fires a software interrupt that marks a detected IP (IRQ_{IP} in Figure 2.14).

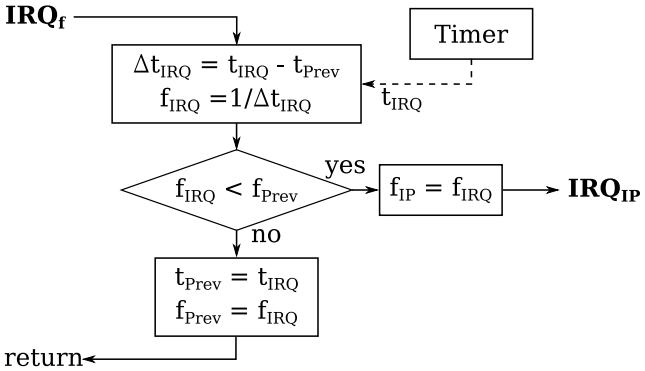

Figure 2.14: Program flow of the interrupt routine for frequency measurement of V_{INT} and IP detection.

The IP marks the point in time, at which the excitation of the harvester has finished. The output frequency of the harvester at IP, f_{IP}, indicates the intensity of the excitation. Thus, it is related to the expectable *power supply duration* (PSD) time t_{PSD}. To determine the relationship between f_{IP}, the system load I_{Load} and t_{PSD}, an autonomous data set generation phase was developed. During this phase, the MCU determines the time between the

IP and power off together with the frequency f_{IP} at the IP. Once the IP has been determined, an internal comparator of the MCU is activated to monitor the supply voltage and compare it against a reference voltage. This reference voltage is slightly higher than the minimum operating voltage of the MCU, thus it is used as threshold for detecting upcoming power losses like in the *voltage-based power loss detection* (V-PLD) method (see also Section 2.1.4). Once the supply voltage falls below the reference voltage, the interrupt routine of the comparator, IRQ_{Comp}, is entered. Inside this routine, the power supply duration time t_{PSD} is determined as the passed time between the preceding IP detection and the current supply voltage drop. This value is grouped together with the load current, I_{Load}, and the frequency at the IP, f_{IP}, and written as data set to the PSD list.

More details about the implementation, data set generation, data preprocessing are described in the related publications [47, 51].

Evaluation and Results

Figure 2.15 gives an overview about the measurements in the data set generation phase on the evaluation hardware once the system is out of cold-start phase. It outlines the coherence of the voltage levels and relevant points in time for a single harvester output burst.

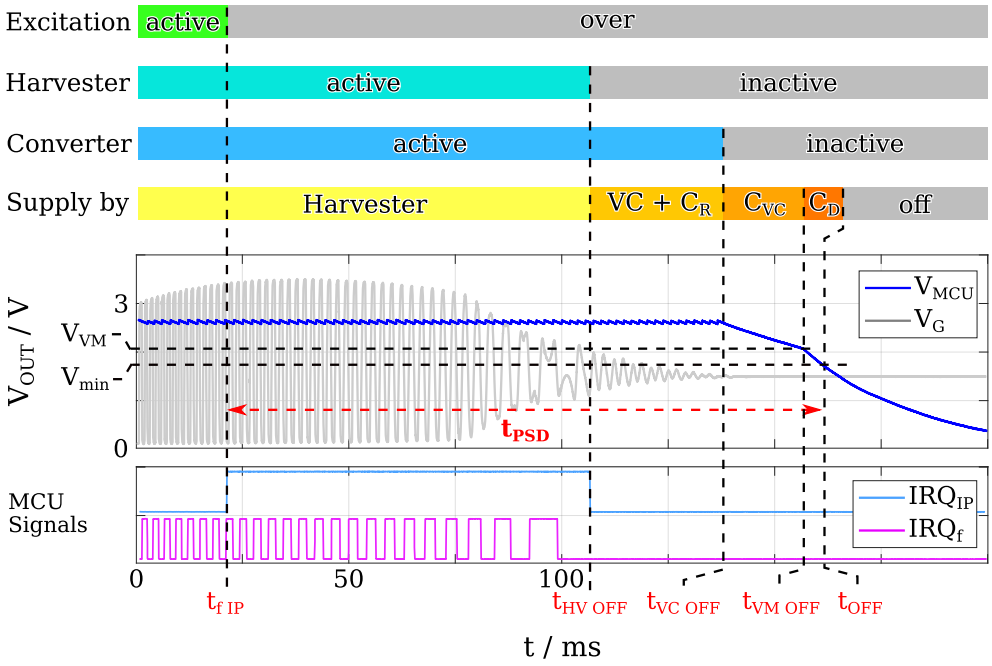

Figure 2.15: Evaluated voltage and time measurements of a single output burst.

Once the harvester output voltage falls below the minimum input voltage of the *voltage converter* (VC) at $t_{HV\ OFF}$, the VC continues providing power due to the remaining load in the capacitors of the rectifier C_R until $t_{VC\ OFF}$. After that, the voltage begins to drop,

as the remaining load of the capacitors C_{VC} of the VC declines, until V_{MCU} is power-gated by the *voltage monitor* (VM) of the cold-start unit at $t_{VM\ OFF}$. For a short time, the system runs solely on the load in the decoupling capacitors C_D, until V_{MCU} falls below the minimum operating voltage V_{min} of the MCU at t_{OFF}.

The signal IRQ_f indicates the pin interrupts of the frequency measurement, based on the harvester output signal V_G. The proposed detection algorithm of Figure 2.14 determines $t_{f\ IP}$ as inflection point of the harvester output frequency, indicated as rising edge in the IRQ_{IP} signal in Figure 2.15.

The data set generation phase has been conducted multiple times to determine the relationship between the intensity of excitation, the system load and the resulting PSD. For this evaluation, one hundred measurements with different excitations have been recorded for each system load setting. This number became apparent to be sufficient for an expressive evaluation of the frequency dependency of the PSD. The voltage threshold V_{th} for the power loss identification of the supply voltage has been set to $2V$ (equal to V_{VM} in Figure 2.15).

The raw measurement results have been used to derivate a *guaranteed minimal power supply duration* (gmPSD) time at the different system loads for specific periods of f_{IP}. More details are described in the related publications [47, 51].

Figure 2.16 shows the results, i.e., the gmPSD time for all system load settings. The *F-PLD* curves are the outcome of a regression analysis based on the raw measurements and represent the functional relation for a continuous f_{IP}. Due to the computational overhead, this relationship has however not been implemented in a functional manner but as *look-up table* (LUT) over a quantized f_{IP}. Using the LUT only incorporates a very low overhead for mapping the currently measured f_{IP} to a specific LUT index.

The results show, how the output frequency of the harvester at the IP, f_{IP}, affects the PSD in each specific system load setting. High intensities of harvester excitation, reflected as high output frequencies lead to a longer PSD than lower intensities. The results also show, that the PSD differs for different system load settings. Higher loads, caused by higher settings of f_{CPU}, lead to a lower PSD at the same f_{IP}.

In order to compare these results with the state-of-the-art method, the achievable results with the V-PLD method at the same decoupling capacity are calculated by Equation 2.1 and added as red-dashed lines. The calculation is based on a threshold voltage of $2.17V$ as used in [62] and a V_{min} of $1.85V$, which is the upper limit of the MCU off voltage according to the data sheet. Most related PLD methods have been designed around a constant decoupling capacity value, as it is present on the employed evaluation hardware. In contrast, our new principle is agnostic to the decoupling capacity. Whereas our evaluation only employs the minimum recommended value of $2.2\mu F$ [6], V-PLD-based approaches require values that are about five [94] to ten times larger [61, 62].

[6] According to the data sheet of the MCU, a minimum decoupling capacitance of $C_{D\ min} = 2.2\mu F$ is recommended.

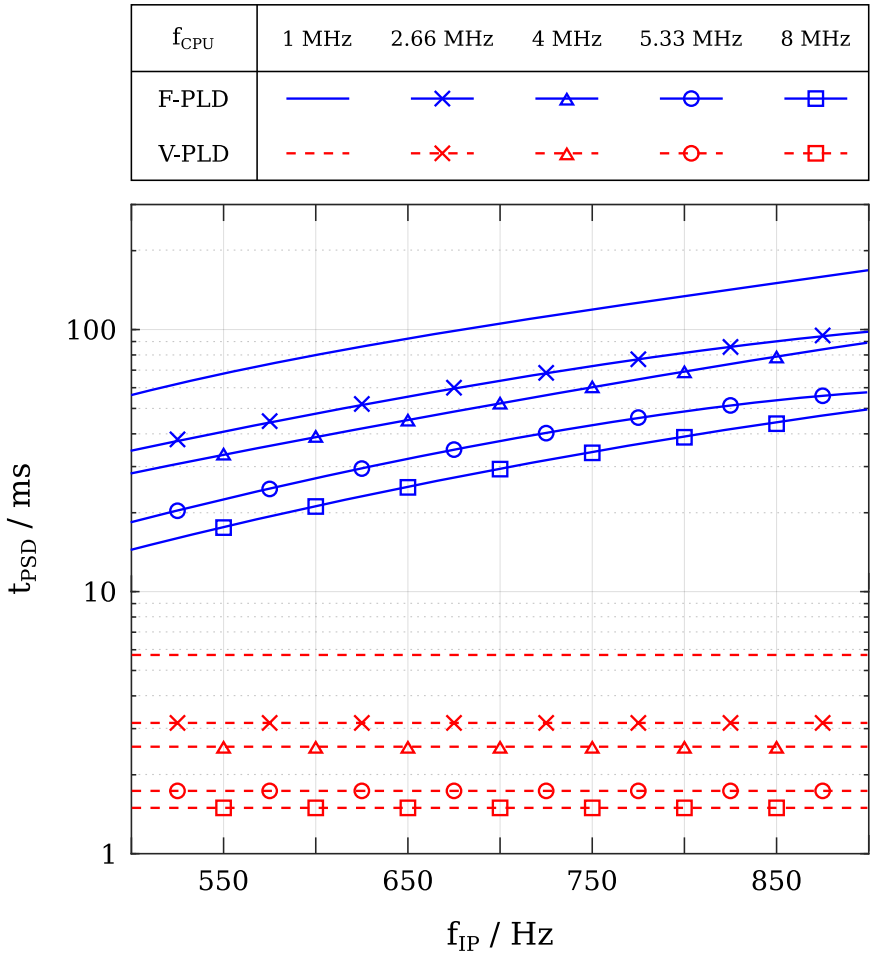

Figure 2.16: gmPSD time for specific system loads at different IP frequencies of the *frequency-based power loss detection* (F-PLD), compared to V-PLD based results with the recommended decoupling capacity.

Figure 2.16 shows that our new F-PLD method is capable of detecting power losses more than 10 times earlier, and thus allows to schedule the software of *Power-Neutral* systems in a *proactive* manner by considering upcoming power losses early on. The dependency of the IP frequency indicates, that the F-PLD method is sensitive to the dynamics of the harvester, whereas the V-PLD-based PSD is constant over the output frequency of the harvester.

To evaluate the cost of our approach, the power overhead is analyzed and compared with the overhead required for V-PLD. The overhead of V-PLD methods is defined by the additional energy that is consumed to monitor the supply voltage. Early works on CP reported relatively high overheads for the V-PLD monitoring [62] and employed energy-hungry comparators [61]. In contrast, the use of custom external monitoring circuits can lead to significantly lower overheads than using MCU-internal peripherals [94]. For this evaluation the V-PLD was implemented with internal peripherals of the MCU: The comparator and voltage reference cause a continuous power overhead of around $55\mu W$.

Regarding our F-PLD approach, the overhead is dominated by the power dissipation in the resistor R_{INT} in the frequency measurement bypass path (Figure 2.13). For a value of $R_{INT} = 150k\Omega$ it was measured at different output frequencies f_G of the harvester. Figure 2.17 shows the power overhead P_{OH} of the V-PLD method (red line) and the F-PLD method (blue line).

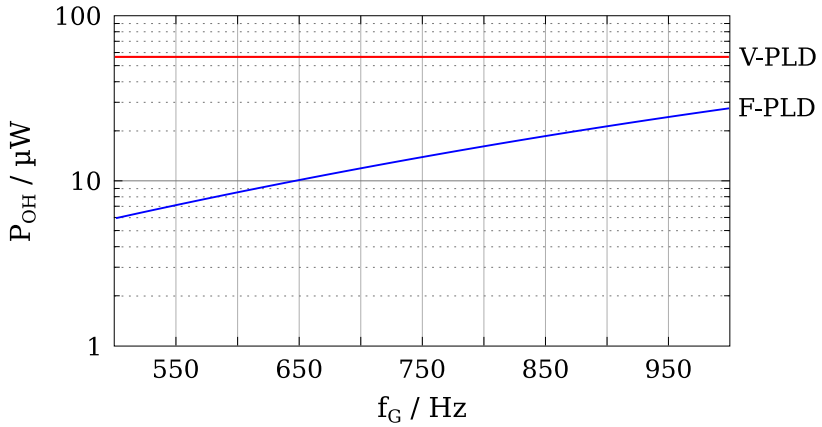

Figure 2.17: Power overhead of the F-PLD compared to the V-PLD method.

It indicates that the overhead of our new F-PLD method is lower than of V-PLD on our evaluation hardware. Since the voltage drop across the resistor depends on the harvester output voltage, which is again related to the frequency, the power overhead depends on the output frequency.

Discussion and Conclusion

The presented principle considers the internal state of kinetic energy harvesters to identify upcoming power losses significantly earlier than state-of-the-art voltage-based PLD approaches. Therefore, the kinetic energy that is mechanically stored in the excited rotating mass is exploited as a predicable source of power. To determine the amount of this energy, which is reflected in the intensity of the excitation of this mass, the output frequency of the harvester is measured. The reported results show, that potential power losses can be detected more than 10 times earlier than state-of-the-art voltage-based approaches. Therefore, this method offers transient computing that is adaptive to the dynamics of the harvester and enables a *proactive* power loss handling without the need of additional decoupling capacity for checkpointing. Whereas state-of-the-art *reactive* concepts must stop the application abruptly for instant backup of the system state, the *proactive* handling can ensure the finalization of atomic operations.

2.2.2 Predictive Checkpointing for Thermoelectric Harvesters[7]

In this section, a new method is presented to detect upcoming power losses before the voltage begins to drop for transiently powered IoT sensor devices with thermal energy harvesting. As explained in the appendix B.2, the output voltage V_{TEG} of a TEG depends on the applied temperature gradient ΔT, the Seebeck constants of the incorporated materials, and the wiring of the internal structure. The presented method utilizes the relationship between the temperature gradient ΔT and the maximum power point output P_{MPP} of TEGs. Specifically, the history of the temperature gradient is monitored and utilized to predict the trend of P_{MPP}. This method furthermore exploits the thermal capacity of a TEG, i.e., the thermal time constant in the correlation between its exciting input quantity and its electrical output. Thermal simulations have shown that the thermal capacity leads to a delayed permeation of the environmental thermal into the internal power-generating structure of a TEG. In contrast to the approach as presented in Section 2.2.1, in which the internal state of the harvester is monitored to predict the future energy trend, the approach in this section monitors the exciting physical quantity of the harvester for power prediction.

In this work, it is distinguished between two temperature gradients: $\Delta T_{surface}$ and $\Delta T_{internal}$. As shown in Figure 2.18, $\Delta T_{surface}$ represents the gradient at the outer surface of the TEG, whereas $\Delta T_{internal}$ represents the gradient at the internal structure of the metallic elements, in which the thermoelectric effect (Seebeck effect) takes place. Thus, both gradients are thermally separated by the physical properties of the surface material and its geometry.

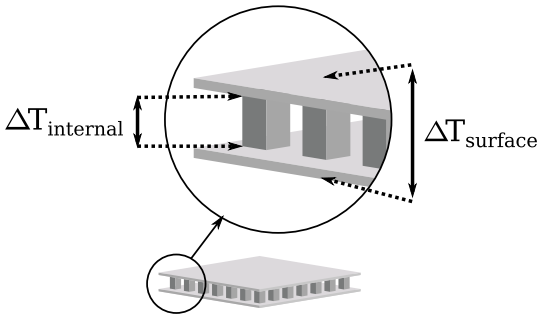

Figure 2.18: The definition of $\Delta T_{surface}$ and $\Delta T_{internal}$ at a TEG.

Methodology and TEG Analysis

The basis of the concept in this work is that the temperature gradient inside the internal structure $\Delta T_{internal}$ is delayed to the gradient at its surfaces $\Delta T_{surface}$ once $\Delta T_{surface}$ starts dropping. Thermal simulations have been performed to investigate the temporal coherence between $\Delta T_{internal}$ and $\Delta T_{surface}$. The outcomes showed, that $\Delta T_{internal}$ follows $\Delta T_{surface}$

[7] This section and the related graphics have already partially been published in [58]

with a delay in the range of several seconds depending on the drop rate of $\Delta T_{surface}$. The delay in general can be explained by the thermal capacity of the TEG and its geometric structure. Thus, the basic idea is that $\Delta T_{internal}$, i.e., the relevant quantity for the power output of the harvester, can be predicted by measuring and analyzing the history of $\Delta T_{surface}$, i.e., the exciting input quantity of the harvester.

To investigate the temporal coherence between $\Delta T_{internal}$ and $\Delta T_{surface}$, an experimental system setup was designed as shown in Figure 2.19.

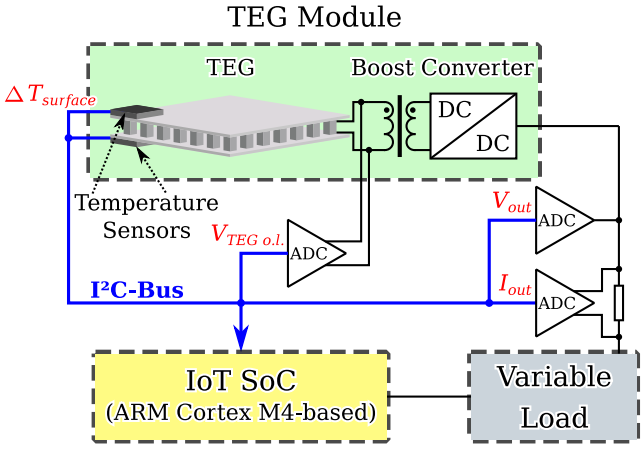

Figure 2.19: Circuitry of the experimental system setup.

This setup incorporates a small TEG module by Matrix [165] with internal boost converter and a common Flash-based IoT SoC. The TEG is equipped with two low power temperature sensors by Maxim Integrated [166]. In order to analyze the relationship between the thermal input ΔT and the electrical power output of the TEG at its maximum power point P_{MPP}, a variable resistive load is employed together with related voltage and current measurement V_{out} and I_{out}. Furthermore, an *analog to digital converter* (ADC) measures the open loop output voltage of the TEG, $V_{TEG\ o.l.}$, which is later on used to derive P_{MPP}.

The temperature gradient that is responsible for the power generation is $\Delta T_{internal}$, which develops itself in the internal structure of the TEG. However, as the internal structure of the TEG is not accessible for temperature sensors, $\Delta T_{internal}$ cannot be measured like $\Delta T_{surface}$. Therefore, a preliminary analysis was conducted under static thermal conditions. Such a static thermal condition is achieved by keeping $\Delta T_{surface}$ constant until $\Delta T_{internal} = \Delta T_{surface}$. This enables a characterization of the TEG module, where specific $\Delta T_{internal}$ values correlate with related P_{MPP} values. Due to the physics of the Seebeck effect, specific $\Delta T_{internal}$ values directly correlate with $V_{TEG\ o.l.}$, i.e., the open loop output voltage of the TEG. Thus, again specific $V_{TEG\ o.l.}$ values correlate with specific P_{MPP} values. This correlation is exploited in the subsequent analysis, as $V_{TEG\ o.l.}$ can be measured faster than P_{MPP} and therefore yields more reliable results under fast changing thermal conditions. In other words, the results of the static analysis, shown in Figure 2.20, demonstrate the coherence between the physical input quantity $\Delta T_{internal}$ and the related output

quantities $V_{TEG\ o.l.}$ and P_{MPP} of the employed TEG module. Figure 2.20 also shows the minimal operating power $P_{op\ min}$ of the SoC that has been employed. Thus, a shortfall below the specific power output level $P_{op\ min}$, i.e., the minimal operating power of the electronic subsystem, is defined as a power loss. Here, $P_{op\ min}$ is set to $0.25mW$, since this is the minimal operating power at the lowest possible clock frequency of $100kHz$ during code execution from Flash at a supply voltage of $2.5V$ [167].

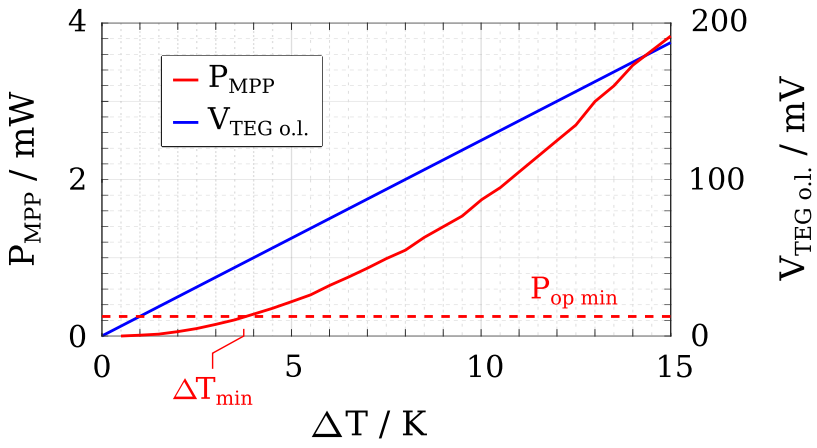

Figure 2.20: Relationship between $V_{TEG\ o.l.}$ and P_{MPP} over ΔT under static thermal conditions.

In order to quantify the time shift between $\Delta T_{surface}$ and $\Delta T_{internal}$ under changing thermal conditions, further measurements and analyses have been conducted. As a starting point for the measurements, a temperature gradient of $\Delta T_{surface} \approx 12$ Kelvin has been applied to the TEG. This value has been chosen as upper limit for thermal energy harvesting applications that scavenge the temperature gradient from environmental sources. After static thermal conditions have been met, i.e., $\Delta T_{surface} = \Delta T_{internal}$, a potential power loss has been emulated by removing the applied temperature gradient and leaving the TEG setup in an environment with room temperature where $\Delta T_{surface}$ immediately starts dropping. Concurrently, the dropping gradient $\Delta T_{surface}$ and the respective values of $V_{TEG\ o.l.}$ have been monitored. The values of $V_{TEG\ o.l.}$ have been used to determine the concurrent $\Delta T_{internal}$ by using the coherence shown in Figure 2.20. Figure 2.21 shows the progression of $\Delta T_{surface}$ and $\Delta T_{internal}$ of one measurement during a dropping temperature gradient over time. Once the gradient begins to drop after the static thermal condition at $t(0)$, a continuous time shift between $\Delta T_{surface}$ and $\Delta T_{internal}$ can be observed, exemplarily displayed as red arrow $\Delta t_{surface-internal}(t_{min})$ at t_{min}. The dashed line represents $\frac{d}{dt}\Delta T_{surface}$, i.e., the time derivative of $\Delta T_{surface}$, which shows a minimum in the beginning at t_{min}. The statistical evaluation of the analyses revealed that after this minimum the time shift directly correlates with the intensity of $\frac{d}{dt}\Delta T_{surface}$. ΔT_{min} represents the minimal temperature gradient to guarantee the minimal operation power $P_{op\ min}$ of the employed SoC (see Figure 2.20). Therefore, a shortfall of $\Delta T_{internal}$ below ΔT_{min} is considered as a power loss event for the experimental system.

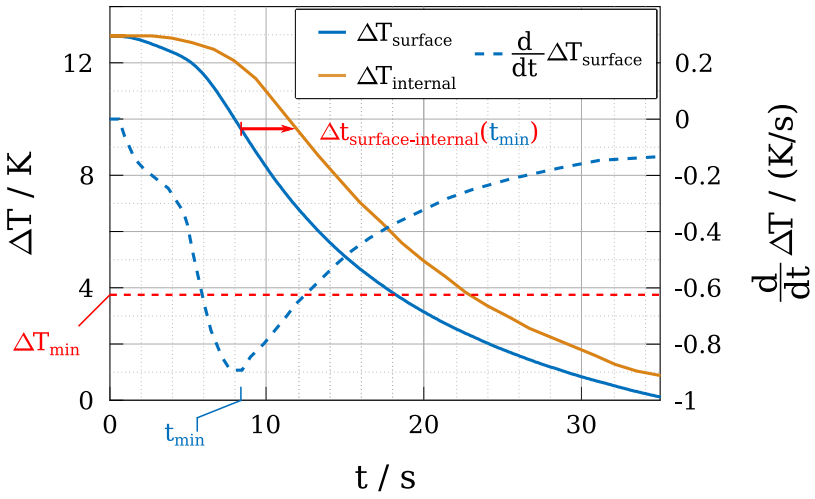

Figure 2.21: Progression of $\Delta T_{surface}$ and $\Delta T_{internal}$ during a power loss emulation together with the time derivative $\frac{d}{dt}\Delta T_{surface}$ of $\Delta T_{surface}$.

Evaluation and Results

To evaluate the correlation between the time shift and the intensity of $\frac{d}{dt}\Delta T_{surface}$, multiple measurements have been performed and analyzed. Figure 2.22 shows the measurement results and how the time shift $\Delta t_{surface-internal}$ correlates with $\frac{d}{dt}\Delta T_{surface}$ after its minimum (represented in Figure 2.21 at t_{min}) has been passed. Whereas $\Delta t_{surface-internal}$ shows a low deviation and clear trend for $\frac{d}{dt}\Delta T_{surface} < -0.2$ Kelvin per second, it spreads increasingly for values of $\frac{d}{dt}\Delta T_{surface}$ closer to zero. Since $\Delta T_{internal}$ has already fallen short of ΔT_{min} at $\frac{d}{dt}\Delta T_{surface} = -0.2\frac{K}{s}$ (as shown in Figure 2.21), the evaluation of $\frac{d}{dt}\Delta T_{surface}$ closer to zero is not relevant for the system setup in this work and won't be for other Flash-based IoT SoCs.

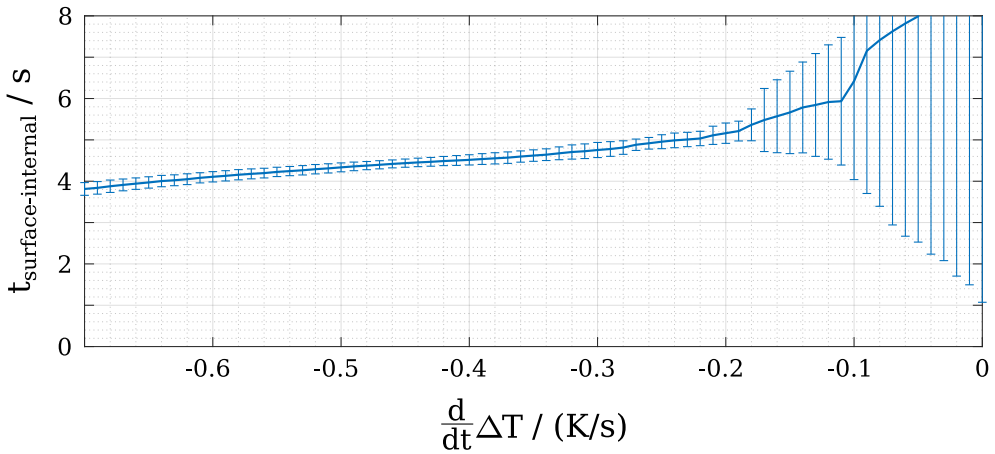

Figure 2.22: Mean value and standard deviation of $\Delta t_{surface-internal}$ with respect to $\frac{d}{dt}\Delta T_{surface}$ after its minimum.

The results in Figure 2.22 indicate that the time shift $\Delta t_{surface-internal}$, i.e., after which $\Delta T_{internal}$ will settle to the current $\Delta T_{surface}$, can be determined based on the value of the time derivative $\frac{d}{dt}\Delta T_{surface}$. In other words, time derivative values below $-0.2\frac{K}{s}$ allow a prediction of the internal temperature gradient of the TEG $\Delta T_{internal}$, based on the historic measurement data of its surface $\Delta T_{surface}$. Therewith, the power output P_{MPP} trend can be predicted by utilizing the coherence between P_{MPP} and ΔT as shown in Figure 2.20 by setting $\Delta T = \Delta T_{internal}$. As an example, the power loss of the experimental system setup at $P_{op\ min} < P_{MPP}$ occurs at around $t = 23s$ in Figure 2.21. By exploiting the relationship of Figure 2.22 together with the historic data of $\Delta T_{surface}$, this power loss can be predicted more than four seconds in advance. Since $\Delta t_{surface-internal}$ exhibits only a low slope over $\frac{d}{dt}\Delta T_{surface}$, it also allows to simply utilize the minimal $\Delta t_{surface-internal}$ value of ≈ 3.5 seconds. This largely reduces the complexity of the temperature data analysis, as only the detection of a drop of $\Delta T_{surface}$ is needed.

Our new method enables an efficient predictive CP in Flash-based IoT edge devices, as the energy-intensive write procedures require a guaranteed power level for a specific amount of time to ensure a reliable CP backup process. The latest point in time, i.e., the point in time where P_{MPP} will drop under the minimal requirements of a successful Flash write procedure, can be accurately predicted with a minimal error to timely trigger the CP backup process.

The power demand of the temperature sensors causes a power overhead for the proposed principle. This is shown in Figure 2.23. The overhead strongly depends on the duty cycle, i.e., the sample frequency f_s, of the temperature measurement.

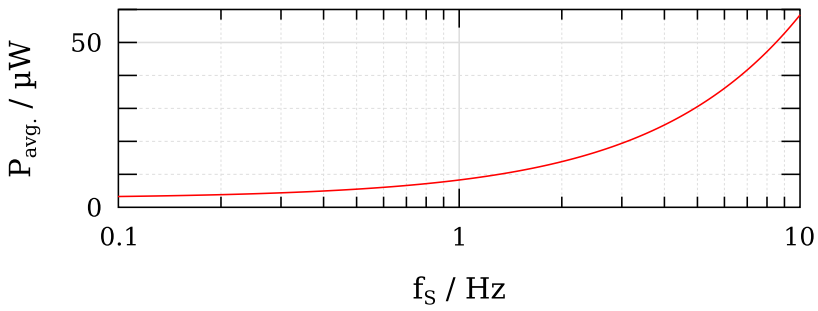

Figure 2.23: The average power overhead $P_{avg.}$ of the temperature sensors depending on the sample rate f_s at a supply voltage of 2.5V.

During normal operation, a sample rate of $f_s = 1Hz$ or less is sufficient to timely detect a dropping of $\Delta T_{surface}$. Thus the average power overhead will be smaller than $10\mu W$, which is less than 4% of the operating power of the SoC at its minimal operating frequency. This percentage decreases even further at higher operating frequencies. Increasing f_s might only be needed during a dropping $\Delta T_{surface}$, in order to determine the drop rate more accurately.

Discussion and Conclusion

It has been demonstrated, that the future trend of the power output becomes predictable by monitoring the exciting quantity of the harvester, i.e., the temperature gradient at the surface of a TEG and exploiting its thermal capacity. The reported results show, that once the temperature gradient at the surface of the TEG decreases faster than 0.2 Kelvin per second, its development inside the power-generating structure is predictable with a standard deviation of down to 200 milliseconds. As the internal temperature gradient is decisive for the intensity of the generated power, critical power levels can be predicted several seconds before they occur. This new method therewith enables to detect potential power losses more than two orders of magnitude earlier than state-of-the-art voltage-based methods (Equation 2.2). To the best of our knowledge, this is the first approach to ensure a reliable checkpointing backup process on Flash-based IoT devices without adding large amounts of capacitance to the system as in [129, 130].

2.2.3 System Cold-Start Optimization

The system cold-start describes the behavior of an embedded system when it starts up without prior any energy in the system. This cold-start phase begins when the energy harvester starts to provide electrical power to the system and ends with the MCU executing the first instruction of the application code. Typical initial conditions of embedded systems prior to the cold-start phase are:

- No power supply
- All capacitors are empty
- All *integrated circuits* (ICs) and their internal oscillators if any, are off. No software code execution
- No data retention in volatile memories

In *Energy-Neutral* or conventional battery-powered embedded systems, the cold-start phase only occurs once in the beginning of the lifetime or once upon a battery replacement cycle. To conserve energy or to react on a temporary weakening power supply, the systems incorporate energy saving modes like idle, sleep or power-down mode. These modes usually come with an optimized tradeoff between low energy demand and fast system wake-up, which allows a calculable behavior of the system that can be controlled by the application software. However, to make use of energy saving modes and fast system wake-ups, a reliable energy buffer that contains a specific residual amount of energy is required to prevent the system from experiencing a complete power loss. Due to the absence of energy buffers, *Power-Neutral* transient computing systems naturally experience cold-starts all the time. In contrast to *Energy-Neutral* or battery-powered systems, they most likely will experience a complete power loss once the power supply is weakening. Figure 2.24 gives an overview about how the different system types will react on a weakening power supply, i.e., when it doesn't provide sufficient power for the active mode of the system.

t_0 depicts the point in time, where the power supply unit starts providing power for the first time. t_{CS} is the cold-start time and t_{weak} the point where the power supply starts weakening and, after a short period of time, returning to normal at t_{ok}.

As emphasized in Figure 2.24, the cold-start energy E_{CS} is usually significantly larger than the wake-up energy $E_{Wake-up}$, E_{CS} depends on many aspects in the circuitry, such as the capacity of the capacitors, the power output of the energy harvester as well as the cold-start timing, undervoltage behavior and minimal operating voltage of specific submodules like the voltage converter. Weak power outputs of the harvester can result in a long cold-start phase, which appears as a long and inconvenient switch-on delay of the embedded system. Energy harvesters with a short power output duration, such as kinetic types, might even be incapable of starting-up the system at all (e.g. $t_{weak} < t_0 + t_{CS}$ in Figure 2.24). Therefore, minimizing the cold-start energy E_{CS}, and therewith the cold-start time t_{CS},

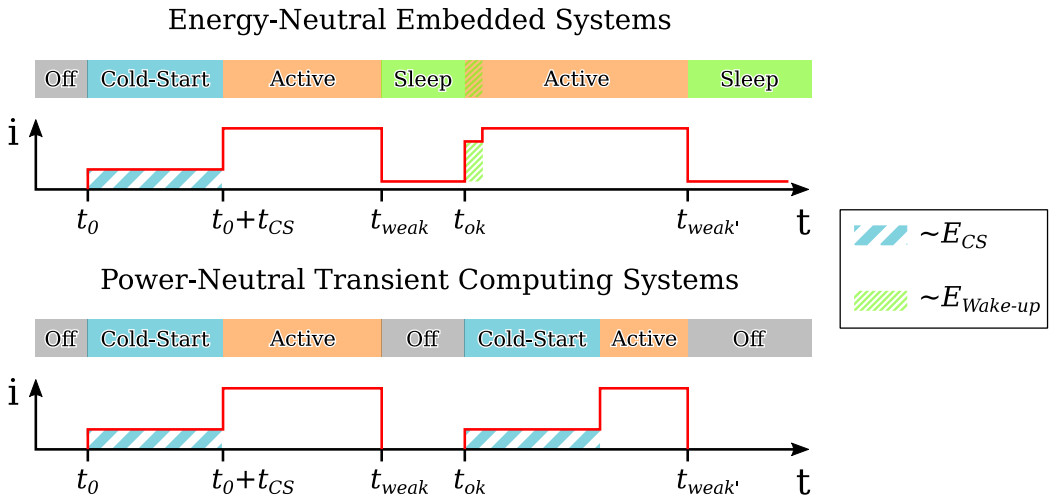

Figure 2.24: The impact of the cold-start phase for *Power-Neutral* in contrast to *Energy-Neutral* systems elucidated by means of the current profile. (For the sake of simplicity, currents are assumed to be constant during a phase.)

is an essential design goal for a sophisticated *Power-Neutral* transient computing system design. For this, however, the challenging conditions prior to a cold-start and the respective constrains must be faced:

- Initially, the entire system behaves as a very low impedance load to the energy harvester due to the empty capacitors.
- Since energy harvesters usually have a weak power output, it is difficult to overcome the initial low system impedance until the capacitors are charged to the minimum operating voltage.
- As soon as the minimum operating voltage of a submodule, i.e., a voltage converter, is reached, its current demand increases abruptly. This can cause the voltage to drop again below the minimum operating voltage and thus lead to dead-lock oscillations.

Surprisingly, there is almost no specific related work about this issue for *Power-Neutral* systems. Existing proposals focus on optimizing specific components in the voltage supply chain, such as the rectification of kinetic harvester output voltage. Here, alternatives to diode-based rectifiers have been proposed like dedicated chopper circuits [98], rectifier concepts with negative voltage converters [96], or rectifiers that switch between active and passive mode, depending on the available power in the system [97]. Then again, Balsamo et al. [94] proposed a promising concept for a start-up circuit to prevent the current demand of MCUs in *Undervoltage Conditions* (UVCs), which has been employed as cold-start unit in the proposed voltage conditioning chain shown in Figure B.2. As far as the voltage conversion is concerned, the LTC3129 chip by Linear Technology has been evaluated to come with better cold-start properties than other DC/DC converters [63], which is why this chip has been used in the proposed hardware in this work shown in Figure B.1. But most

proposed system concepts consider their system to be out of the cold-start phase or even ignore this state. For instance, popular CP-related publications employed laboratory power supplies with pulsed output to simulate power losses and returns [61, 62]. Due to the low-impedance power output of these power supplies, the related cold-start phases had been very short and became irrelevant for their system designs.

Composition of the System Cold-Start Energy

In order to reduce the cold-start energy overhead, its sources must be identified and evaluated. Figure 2.25 shows typical sources of cold-start energy in the submodules of the voltage conditioning chain and the MCU in a kinetic energy harvesting-powered *Power-Neutral* embedded system.

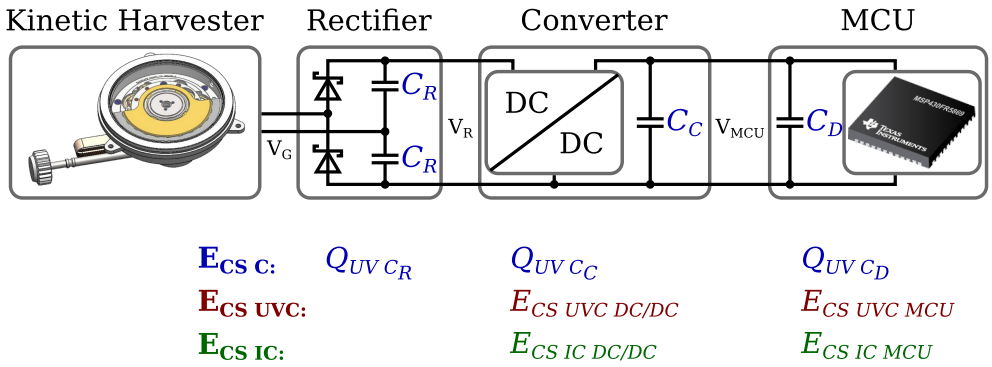

Figure 2.25: Typical sources of cold-start overhead energy in the submodules of a *Power-Neutral* transient computing systems. For further details about the shown hardware components see Chapter B.1.

$\mathbf{E_{CS\ C}}$ represents the energy overhead for charging the capacitors of their related submodules/*integrated circuits* (ICs) during their UVC, i.e., from zero volt in the beginning at t_0 until their minimal operating voltage is reached at $t_{V_{IC\ min}}$. For instance, $Q_{UV\ C_D}$ represents the charge that is required until C_D is charged to the minimal operating voltage of the MCU $V_{MCU\ min}$:

$$Q_{UV\ C_D} = \int_{t_0}^{t_{V_{MCU\ min}}} i_{C_D}(t)dt \; or \; rather \; Q_{UV\ C_D} = V_{MCU\ min} * C_D \tag{2.5}$$

As the capacitors are initially empty, they represent a very low impedance load in the beginning, which is challenging for energy harvesters with weak power outputs.

$\mathbf{E_{CS\ UVC}}$ represents the energy overhead that is drawn by the submodules/ICs until their minimal operating voltage is reached, i.e., once they are out of their UVC:

- $E_{CS\ UVC\ DC/DC}$ represents the energy demand of the DC/DC converter during UVC.
- $E_{CS\ UVC\ MCU}$ represents the energy demand of the MCU during UVC.

$\mathbf{E_{CS\ IC}}$ represents the energy overhead that is drawn by the submodules/ICs during their start-up phase, i.e., after they are out of their UVC until they normally operate:

- $E_{CS\ IC\ DC/DC}$ represents the energy demand of the DC/DC converter during cold-start period (ramp-up of the switching unit etc.) until the output voltage is supplied.
- $E_{CS\ IC\ MCU}$ represents the energy demand of the MCU during cold-start period (ramp-up of the oscillator, execution of the initialization code, in which stack, RAM, etc. are initialized) until the first execution of application code.

Optimization Strategy

In order to reduce the discussed sources of cold-start energy, the respective constrains and interdependencies must be considered:

- *Reducing* $E_{CS\ C}$:
  - The manufacturer of the submodule/IC usually recommends a minimum value for proper operation.
  - Small values can lead to instabilities as sudden current draws directly lead to voltage drops, which might result in a dead-lock of the submodules due to an oscillation between start-up and entering UVC again.
- *Reducing* $E_{CS\ UVC}$:
  - There is almost no room for optimization, as the behavior in UVC is determined by the internal circuitry of the submodule/IC and usually not specifically defined by the manufacturer.
- *Reducing* $E_{CS\ IC}$:
  - There is almost no room for optimization in nonprogrammable submodules/ICs, as their internal behavior is fixed.
  - Regarding submodules/ICs that are programmable by software or tunable by external circuitry, at least limited room for optimization is provided.

Due to these different constrains, the optimization strategy varies for the different sources of cold-start energy:

Reducing $E_{CS\ C}$

According to Equation 2.5, the required charge to exceed the UVC proportionally depends on the size of the respective capacitor. Therefore, minimizing $E_{CS\ C}$ requires to minimize the capacity. As far as the values of C_C and C_D are concerned, the manufacturers of the respective submodules/ICs define minimum values that must not be deceeded in order to

ensure a reliable operation. Regarding C_R, there is more room for optimization. Low values lead to a fast rise of rectified voltage but in return to a fast drop once a load is applied. High values lead to the opposite behavior. Due to the various interdependencies in the other components in the voltage supply chain, the optimal value of C_R will be determined empirically.

Reducing $E_{CS\ UVC}$

As the behavior of the submodules/ICs in UVC is usually not specifically defined by the manufacturer, the best strategy is to overcome the UVC as fast as possible in order to enter the valid operating range. Therefore, the supply voltage must ramp-up as fast as possible. However, as soon as the voltage reaches the minimum operating voltage, the respective submodules/ICs will most likely induce a sudden current draw. Consequently, the voltage must be stable and loadable at this time to prevent that these current draws cause the voltage to drop towards UVC again. Thus, the strategy for minimizing $E_{CS\ UVC\ DC/DC}$ and $E_{CS\ UVC\ MCU}$ is to supply a stable voltage as fast as possible.

Reducing $E_{CS\ IC}$

The operation of the employed DC/DC converter is tunable by its external circuitry, which can be used to influence $E_{CS\ IC\ DC/DC}$ in a certain range under given conditions. For instance its start-up voltage value $V_{DC/DC\ RUN}$ can be adjusted by an external voltage divider, i.e., the minimal voltage value, at which the internal switching regulator starts-up. This start-up is coupled with a significant increase in current draw, which is why the choice of this value is directly coupled with the current load of the voltage output of the rectifier. Therefore the choice of $V_{DC/DC\ RUN}$ directly controls an essential tradeoff of the voltage conditioning circuitry: Low values are accomplished faster in the rectifier, which leads to short cold-start times t_{CS}, but in return result in a less reliable start-up procedure, as the related current draw more likely ends in an UVC again. High values cause the opposite behavior. The strategy is to empirically adjust this value and investigate the overall system behavior for different values of C_R.

As it is for $E_{CS\ IC\ MCU}$, there is only limited room as the ramp-up of the default oscillator of the employed MCU cannot be changed. Any customization of the behavior is only possible in software, and therewith once the oscillator is running.

Summary

The discussed optimization strategies can be summarized as follows:

- Degrees of freedom in the design space:
  - Varying the value of C_R, in order to empirically figure out the best tradeoff between fast voltage ramp-up and appropriate current loadability.
  - Varying the start-up voltage level of the DC/DC converter $V_{DC/DC\ RUN}$, in order to figure out the best tradeoff between cold-start time and start-up stability.

  - Reducing the cold-start time of the MCU and minimizing its initial current draw by software as fast as possible.

- Metrics for evaluating the results:

  - The total cold-start time t_{CS} of the entire system, i.e., the point in time, from which the harvester starts producing power until the first instruction is executed in the MCU.
  - The total achievable execution time for the MCU t_{exe} at one single power output burst of the energy harvester.
  - The impact of the circuitry on the degree of power neutrality, i.e., the overall system efficiency.

Evaluation and Results

As lowest excitation of the micro-scale kinetic harvester MGS26.4, a half turn of its oscillating mass has been considered as shown in Figure 2.12 in the second last section. The related power output duration is in the range of 60 to 120 milliseconds for the first voltage burst. One or two subsequent bursts usually follow with a delay of 100 to 200 milliseconds due to the internal spring barrel (see Figure 2.12). This has been used as basis for the results to evaluate our methods at the most challenging conditions.

However the first analyses showed, that the employed voltage conditioning circuit was not able to supply the initial current of the MCU at the harvester excitation conditions described above. The internal voltage supervisor circuit of the MSP430FR MCU has a hysteresis of $V_{hyst\ VS} = 100mV$. But the voltage drop due to the initial current draw right after $V_{MCU\ min}$ was reached was quickly more than $V_{hyst\ VS}$, so that the MCU ran into UVC again.

Therefore, a specific power management strategy was implemented for the MCU, which fulfills the objective of the optimization strategy to minimize the current draw as fast as possible. This power management is twofold: First, the initial start-up code, where the MSP430FR initializes its RAM after a power-on reset, has been skipped by creating an implementation of the function `int _system_pre_init (void)` that returns zero. This reduced the start-up procedure such that the first user code instruction was reached within a couple milliseconds after $V_{MCU\ min}$ was reached. Second, the first instructions of the user code immediately switched to the internal very low power oscillator `VLOCLK`), which reduced the MCU current to about $I_{MCU} = 50\mu A$. To further relieve pressure on the voltage conditioning circuit, the MCU was set to power-down mode and an external ultra-low power voltage monitor was configured as wake-up source, once the supply voltage is in a safe range (see Algorithm 2.1). This strategy managed to configure the MCU to a software-controllable state, before the supply voltage dropped by the critical value of $V_{hyst\ VS} = 100mV$.

Algorithm 2.1: The first user code instruction in `main()` to minimize the current demand right after start-up.

```
select VLOCLK as main clock source
enable interrupt for external voltage monitor
if voltage below threshold then
    enter power-down
continue here when voltage level ok
```

In the next step, as discussed in the optimization strategy, was to empirically optimize the voltage conditioning circuit parameters C_R and $V_{DC/DC\ RUN}$. Figure 2.26 shows how these parameters affect the evaluation metrics of the system t_{exe} and t_{CS}. Since the power output duration shows a relatively high variance at the same excitation — which is due to its internal construction, especially the spring barrel (see also Section 2.2.1) — the results are quantized in a specific range.

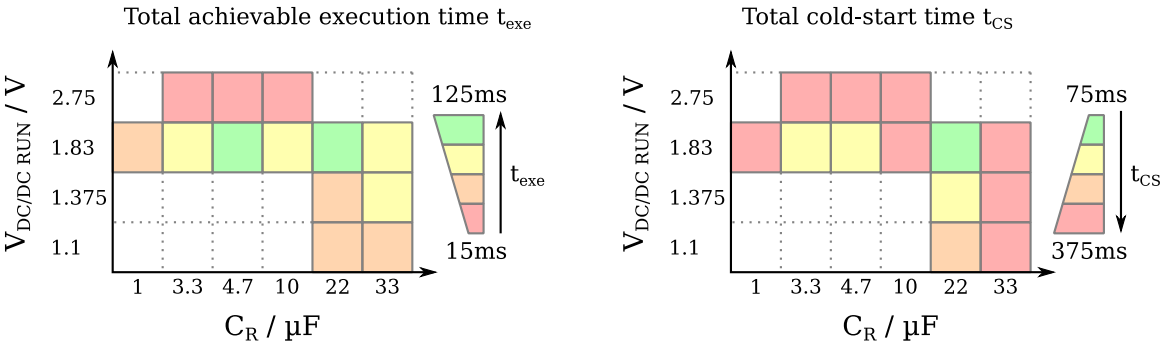

Figure 2.26: Results of the cold-start optimization with respect to execution time and cold-start time (empty fields: No operation possible).

The results show that low start-up voltages $V_{DC/DC\ RUN}$ are only possible with larger capacitors C_R, which has been expected before. However, high start-up voltages result in long cold-start times, even with low capacitors. A middle to high start-up voltage $V_{DC/DC\ RUN}$ combined with medium-size capacitors C_R can be seen as optimal constellation in the proposed system setup. The reason the results partly do not linearly scale with the parameter values, is most likely the different technologies that have been employed for C_R: electrolytic for $1, 22$ and $33\mu F$, foil for $3.3\mu F$, tantalum for $4.7\mu F$ and ceramic for $10\mu F$. Since the option $C_R = 10\mu F$ leads to worse results than its neighbors, the results also indicate that ceramic capacitors might be not a good choice for the rectifier.

Basically, the two parameters evaluated above also define the working point of the kinetic energy harvester, i.e., they define how intensively the harvester is loaded. Due to the physical laws of electromagnetic generators, loads on the harvester cause a breaking action of the generator movement, which in turn reduce its output voltage and frequency. The optimal power point between open loop load and shortcut load is defined as the MPP. Since the MPP depends on different conditions such as the current excitation intensity, it is very

difficult to model. Therefore, this effect has only been evaluated qualitatively based on the output frequency compared to the open-loop frequency.

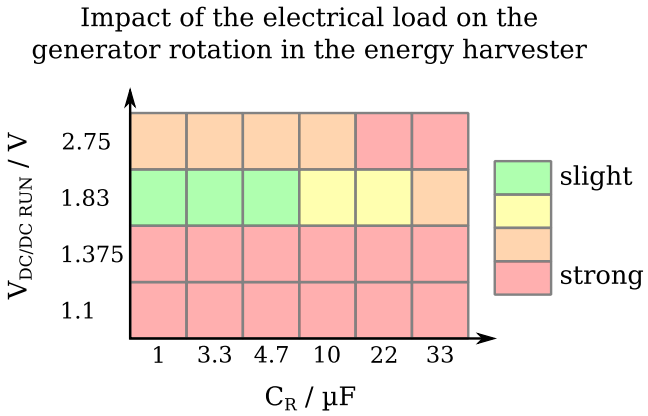

Figure 2.27: Qualitative results of the degree of power neutrality, depending on the design parameters in the voltage conditioning circuit.

Figure 2.27 shows the results, where the conditions *yellow* to *orange* (between the extreme conditions *slight* and *strong*) can be seen as the best conditions, or in other words, as the conditions closest to MPP. Thus, low start-up voltages of the DC/DC converter lead to a too severe burden on the harvester, high values of C_R as well.

Discussion and Conclusion

Related work proposed high capacitors to ensure a stable system start-up despite a weak harvester output, which however led to very long cold-start times [63]. Therefore, the presented approach to minimize the cold-start of *Power-Neutral* systems by minimizing the capacity in the voltage conditioning chain, follows a converse approach to related work. To manage a stable system start-up despite the small capacity, the presented method proposes to increase the start-up voltage of the DC/DC boost converter. This combination led to significant reduction of the cold-start behavior and therewith represents a new design strategy for an MSP430FR-based hardware that ensures a very short cold-start phase to allow a reliable system start-up even from short output duration energy harvesters.

To the best of our knowledge, this is the first work that intensively evaluates the cold-start behavior of *Power-Neutral* embedded transient computing systems. Related work who employed the same energy harvester and the same DC/DC boost converter reported a cold-start time of 2.8 seconds, which required multiple excitations of the harvester [63]. This work in contrast, shows that a cold-start time of less than 75 milliseconds is possible with just one excitation of the harvester.

Chapter 3

Wearable Wireless Sensor Networks for Professional Sport Performance Analysis

$$\vec{\nabla} \times \vec{E} = -\frac{\partial}{\partial t}\vec{B}$$

James Clerk Maxwell

The previous chapter discussed autonomous embedded systems, i.e., battery-free devices that are powered by energy harvesting. This chapter, in contrast, discusses mobile embedded systems as indicated in Table 3.1. More specifically, it addresses energy-specific design challenges of portable devices that are powered by rechargeable batteries.

Autonomous Embedded Systems	**Mobile Embedded Systems**	Stationary Embedded Systems

Table 3.1: This chapter addresses mobile embedded systems. More specifically portable and wearable devices that are powered by rechargeable batteries.

Rechargeable batteries are the usual power supply for portable electronic devices like wearable *wireless sensor networks* (WSNs). The technical requirements depend on the intended application and use case. In professional sports for example, wearable WSNs became very popular as they enable in-field performance analysis in contrast to conventional laboratory-specific measurement equipment like camera sub systems. However, this area of application places high requirements on the overall dimension and weight of the devices, the sensor accuracy and performance, wireless connectivity and battery lifetime. Thus the microelectronic design challenges range from ensuring a high system performance despite the limited energy budget of the battery up to achieving a precise wireless synchronization despite bandwidth-limited low power communication technologies like *Bluetooth Low Energy* (BLE).

This chapter is structured as follows: Section 3.1 covers the technical background, related work and the design challenges of wearable WSNs in professional sports and wireless synchronization principles. Section 3.2 presents the new scientific contributions for ultra-low power system design of wireless sensor networks and proposes new synchronization methods.

3.1 Technical Background, Scientific Challenges and Related Work

3.1.1 Sports Performance Analysis using Wearable Sensor Devices

At the time of this research, commercially available trainer assistance systems came with various disadvantages like high cost and complex setup and calibration procedures. Furthermore, the athlete was required to perform in a preset laboratory environment, making in-field performance nearly impossible.

As a solution, emerging wearable devices of the consumer market, such as fitness trackers, have been applied to professional sports. Together with wireless data transmission, these sensor devices provide online-monitoring possibilities that make real-time performance analysis available to coaches and athletes. In contrast to laboratory measurement setups, they furthermore allow a flexible performance analysis for in-field training.

However, most of these devices have shown to deliver imprecise and unreliable results in the field of professional sports. Even though they provide basic performance measurements such as step counting, they lack of precise analysis and detection of specific disciplines. Therefore, new measurement methods and analysis technologies, dedicated to the requirements of professional sports analysis, have been released within the last decade. Technologies, which deliver more detailed and more precise information about the sports performance of athletes than wearable consumer devices. Such analysis technologies arose manifoldly, e.g., common sports equipment like clubs or shoes that have been equipped with movement sensors or dedicated environments with integrated sensors for movement analysis like contact mats. The employed sensor technologies provide the coaches with biomechanical, physiological, kinetic and kinematic data of an athlete's performance. This data acquisition facilitates highly reliable data analysis-based conclusions about fatigue, training progress and therewith training effectiveness [168, 169].

With the *micro electromechanical system* (MEMS) technology, movement sensors such as *Inertial Measurement Units* (IMUs) became very popular for mobile applications. Therefore, electronic sensor devices became more and more common in sports performance monitoring for in-field usage [64, 65].

IMUs have shown to deliver reliable analysis results for various sports disciplines such as gait and running [170]. Furthermore, IMU-based flight time and jump height determination in counter movements was shown to highly correlate with the concurrently employed optometric reference system [171]. For Drop Jumps, which is a specific discipline for determining the reactive strength index of an athlete, IMU-based parameter determination was shown to highly correlate with a force plate reference system [172]. IMUs are also capable for high dynamic sports movement analysis, but large measurement ranges required due to the high acceleration peaks [173, 174].

Typically, IMUs are a component of a microelectronic sub system that also comes with processing unit and memory. Such systems became popular as senor nodes, i.e., small and lightweight electronic measurement systems that can be worn on the body. To provide maximum flexibility, existing devices are mostly battery-powered. However, this restricts the computational performance of the sensor node's processing unit, since besides the dimensional restriction the energy budget is restricted as well.

Due to the limitations in processing power and energy budget, the implementation of analysis algorithms that run embedded on the sensor node places great challenges and intensive development effort [69]. Therefore, most sports sensor devices simple act as a raw sensor data collector that only write the data to a local memory but do not process the data to compute the desired sports performance parameters. Data analysis has to be done after the raw data recording by using more powerful computing devices like desktop computers that run statistical analysis tools like Matlab [175]. This post-hoc data analysis approach, however, limits the potential benefits of nearly all commercial wearable sports sensor devices, as no real-time feedback can be provided to the athlete during its performance. Even in research, only a few systems provide an on board extraction of basic sports performance parameter based on the sensor data [176].

3.1.2 Time Synchronization for Wireless Sensor Networks

There are situations in professional sports that require a monitoring of the athlete's movement at different locations on the body. Such as specific disciplines that require distributed measurements to, e.g., detect intra-body variations in performance. Thus, such disciplines need multiple instances of sensor nodes for a proper performance analysis. Measurement setups with multiple sensor node instances place new challenges on the system design, such as the requirement for a precise time synchronization between all devices to properly interlink the sensor data of each sensor node. Most commercially available systems follow a wired approach, meaning that all sensors are interconnected using wires. This, however, significantly lowers the flexibility and usability of such systems. Therefore, the only acceptable solution is a wireless time synchronization protocol that comes with sufficient synchronization accuracy and low energy demand.

The conventional *Network Time Protocol* (NTP) was one of the first protocols used for time synchronization in IP-based networks. Since then, several concepts have been proposed more dedicated for wireless networks on resource-constrained embedded devices. *Reference Broadcast Synchronization* (RBS) [177], e.g., employed a reference node that frequently broadcasts timestamps to all nodes in the network. Each node receiving this timestamps can synchronize its internal clock to the given reference. An approach that omits the reference node of RBS was published as *Timing-sync Protocol for Sensor Networks* (TPSN) [178]. TPSN forms a spanning tree of all participating nodes in the network and performs pairwise synchronization along this tree. To reduce the resynchronization

overhead of TPSN, which is required to compensate local clock drifts, the *Flooding Time Synchronization Protocol* (FTSP) [179] was proposed. Therefore, FTSP incorporates a clock drift compensation mechanism.

However, despite their good performance in synchronization accuracy, there are critical drawbacks in all discussed methods. Such as that they are not specifically optimized for a low power operation, which is a critical issue for wearable wireless sensor networks. This becomes particularly apparent by the fast rise of the ultra-low power BLE technology. Furthermore, all mentioned methods, including the *Gradient Time Synchronization Protocol* (GTSP) [180], are based on low layer (e.g. physical layer) timestamps for their timing information. Whereas physical layer timestamps are easily to implement for proprietary wireless protocols, higher-level protocols like BLE usually hide the physical layer activity to the application. The reliability and low power requirements of BLE are based on precise link layer timing. Therefore, no access is provided to the software developer, making all mentioned synchronization protocols theoretically unfeasible for BLE technology. Various approaches of related work either suffer from this fact or came up with proprietary solutions.

Ghoshdastider et al. [70] conducted an experimental evaluation of *Pairwise-Broadcast Synchronization* (PBS) [181] with BLE. Due to the lack of physical layer access, their time synchronization showed an error of $45.071ms$.

Utilizing the acknowledge mechanism of BLE, Bideaux et al. [71] achieved a synchronization latency of $7.51ms$, which is about the chosen connection interval parameter of $7.5ms$. The standard deviation was reported to be $0.41ms$. Utilizing the timing information of the connected event occurrence after a connect command led to an average latency of 39.92μs with a standard deviation of $14.19\mu s$.

Better results were reported with CheepSync [72, 73], a time synchronization service for BLE Broadcasters, in which a custom BLE protocol stack was implemented to provide direct low layer access. Based on hardware timestamps on the transmitter and the receiver side, an accuracy of $10\mu s$ was achieved.

The drawbacks of related work is that either the synchronization error is in the range of milliseconds — which is too high for professional sports performance analysis — or that a custom BLE protocol stack is required. The latter leads to compatibility issues with the official BLE standard and threatens the interoperability among different BLE-enabled devices and communication parties.

As a solution, a dedicated synchronization technique for BLE is presented in this work that utilizes a standard BLE protocol stack and provides a synchronization accuracy in the range of microseconds for two interconnected devices. Furthermore, a proprietary low power wireless *Body Area Network* (BAN) is designed in this work, which supports up to ten distributed units with a sub-millisecond synchronization accuracy among all connected devices.

3.2 New Contributions

This section presents the new scientific contributions of this thesis in the research field of mobile embedded systems, powered by rechargeable batteries. It specifically addresses portable and wearable WSNs in the application area of professional sport performance analysis. The contributions are divides into two topics:

- Section 3.2.1 presents new methodologies for the design of wearable sensor devices. It specifically addresses the design challenges between the goal of a performant and precise movement sensing with real-time analysis capability under the strict constrain of a limited energy budget. A wearable, low power hardware platform design is presented that overcomes the limitation of state-of-the-art wearable sports analysis devices in terms of battery lifetime, real-time performance feedback and in-field usage. This is part of a multidisciplinary project, in which this work covers the system design of the wearable hardware platform. Further contributions are the design of the detection algorithms for the real-time extraction of sports performance parameters [88] and the real-time firmware design and implementation of the detection algorithms[1].
- Section 3.2.2 addresses the design challenges for low power wireless synchronization, which is inevitable for WSNs in use cases that require multiple and distributed sensor devices. It presents two new approaches that exhibit a better synchronization accuracy compared to synchronization principles proposed by related work. The first is based on BLE technology and intended to synchronize two interconnected units with each other, the second presents a specific wireless BAN architecture that allows for multiple units to be synchronized.

[1] Discussed in the upcoming dissertation and the publications of Sebastian Wille

3.2.1 Wearable Sensor Platform

A suitable wearable sensor device must be capable of accurately determining biometrical and kinematic key analysis parameters of the athlete that are typical for such disciplines like flight and ground contact times. Related work has shown, that these kinematic parameters can be analyzed using IMUs [66–68, 170, 182, 183]. Therefore, an electronic subsystem has been designed that employs an IMU as central sensing element.

To develop a sensor system that is suited for flexible in-field performance analysis, the system must be designed as a battery-powered wearable device that can be attached to the athlete's body during his sports performance. Furthermore, the system must be small and lightweight to not affect the movement of the athlete [184, 185]. To allow fast feedback in training situations, the system must provide a real-time event detection and a wireless communication with a human interface host device like a smartphone that directly forwards the analysis data to a coach.

The development of the sensor device has been done in cooperation with the department of sports science at the Technische Universität Dortmund [186] in the context of the SpoSeNs [187] project. Due to the multidisciplinary nature of this project, a dedicated fast prototyping approach was followed to process the different work packages in parallel. Therefore, an evaluation kit (Figure 3.1) was developed that only employs existing hardware modules such as evaluation boards of the IMU sensor and *Microcontroller Unit* (MCU). This kit had been developed in a short period of time and emulates the later wearable sensor device. In order to have the same dimensional properties, the IMU sensor board was placed in the same plastic case as scheduled for the final sensor device. The sensor interface has been forwarded by a cable to a host box that incorporated the batteries, MCU and an SD-Card slot for recording the raw sensor data on.

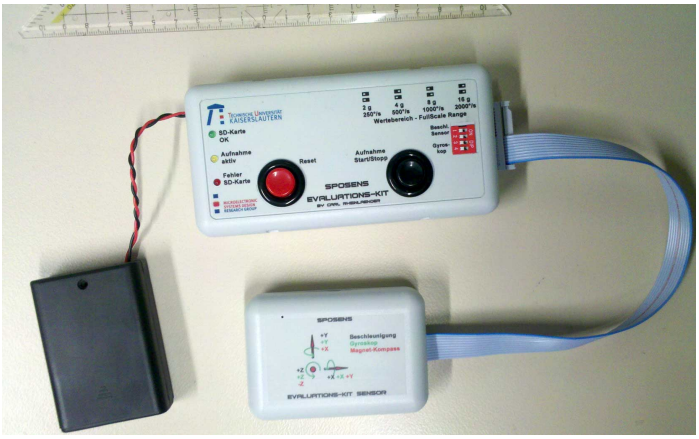

Figure 3.1: SpoSeNs evaluation kit used for early emulation of the wearable sensor device.

This fast prototyping approach enabled an early development of the detection algorithms for the real-time extraction of sports performance parameters [88], together with all related tasks shown in Figure 3.2, before the wearable sensor platform design has been finished.

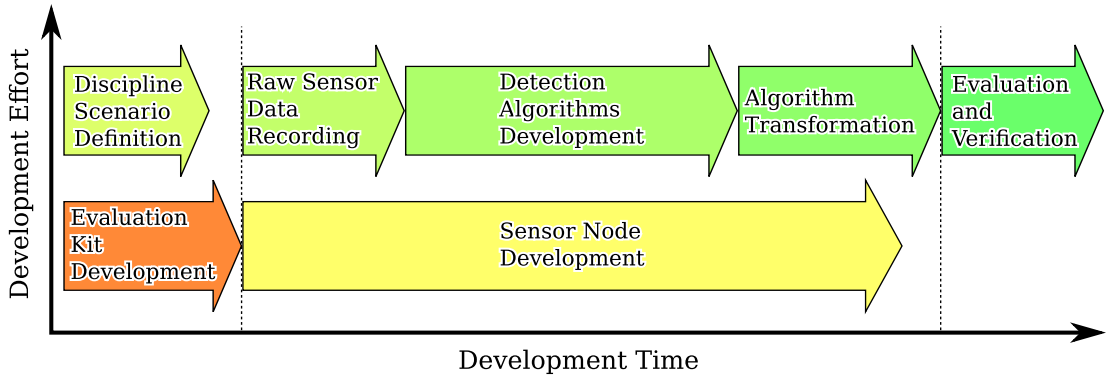

Figure 3.2: The fast prototyping approach enables a reduced development time due to additional development of an evaluation kit. Refer to [88] for more details about the work tasks displayed in the top region.

System Design

The device comes with all the features of state-of-the-art wearable sensor devices for professional sports for fast and precise IMU sensor data recording. In contrast to related work, which is mostly limited to post-hoc data analysis as discussed in Section 3.1.1, the present system design provides additional features for in-field real-time performance feedback to the athlete or coach. Therefore, it is equipped with BLE, a wireless technology that enables low power connectivity with third-party end devices like smartphones, which are used to visualize the computed sports performance parameters to the coach just in time. The software framework is based on an interrupt-based event management with direct memory access to a fast and high capacity NOR-Flash memory. This ensures a fast and precise internal timing that provides a dedicated processing window after each sensor data sample for the real-time execution of detection algorithms[2]. The interrupt-based architecture furthermore guarantees a low power operation of the hardware, as the MCU rests in idle mode most of the time. Once a new sensor sample by the IMU is available, the MCU activated for a short period of time to obtain this data and to execute the next iteration step of the detection algorithm. Therefore, the tradeoff between low power operation and high performance computation mainly depends on the implementation strategy and computational load of the respective detection algorithm.

Figure 3.3 shows the structure of the hardware, which is based on a high-speed ARM Cortex M3 MCU and further peripherals such as a sub-GHz, long-range connectivity technology for future use as well as an air pressure sensor to detect relative height differences in the field.

The small and lightweight sensor device is powered by a lithium polymer battery, that is rechargeable via a micro-*Universal Serial Bus* (USB) port. This battery allows more than 24 hours of active run-time and — by using the micro-SD card — can record nearly unlimited

[2]More details can be found in the upcoming dissertation and the publications of Sebastian Wille

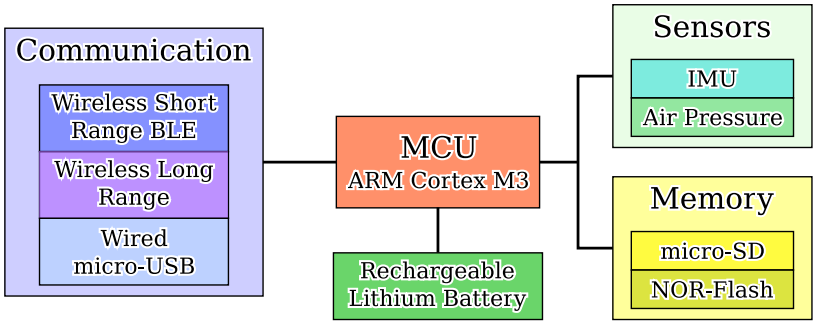

Figure 3.3: Hardware structure of the wearable sensor device platform.

amounts of raw IMU sensor data, i.e., three-axial accelerometer and three-axial gyroscope sensor data, at a sample rate of $1kHz$. This raw data recording feature is intended for the evaluation and development of new detection algorithms for analyzing upcoming sports disciplines. Figure 3.4 shows the final version of the wearable sensor device *printed circuit board* (PCB) and the enclosure[3].

Figure 3.4: Hardware and enclosure of the final wearable sensor device *SpoSeNs Node*.

Specific Applications and Performance Results

Several scientific studies were conducted by utilizing the *SpoSeNs Node* as wearable inertial measurement and real-time analysis device. These studies covered several disciplines of high performance sports, such as:

- *Drop Jump*
- *Sprint*
- *Long Sprint*
- Several pilot studies

Drop Jumps and Sprint are spatially limited disciplines, therefore they have been chosen to additional employ conventional measurement setups as reference systems. The results have

[3]More details about the hardware and firmware can be found in the master thesis of the author [188]

been used to optimize the performance of the *SpoSeNs Node* and to validate the temporal detection mechanisms and their quantitative accuracy. Based on these proofs of reliability, the *SpoSeNs Node* was employed for the other disciplines that where performed in-field.

Drop Jumps are a well-known discipline to evaluate the strength of specific musculature of an athlete. Characteristic quantities for this evaluation are the duration of ground contacts, i.e., the stance duration as well as the flight time after jumping off again. As reference system a force plate (AMTI BP 600400) was employed, on which the athletes performed the drop jumps. The results revealed a detection rate of 94% and a mean absolute error of $3.4 + / - 3.0ms$ for stance duration and $4.9 + / - 3.9ms$ for the flight time using the *SpoSeNs Node* [35].

For the *Sprint* discipline a light barrier system (OptojumpNext) was employed as reference system. The results revealed a step detection rate of 95.7%, whereas for half of the athletes all steps were detected correctly. The stance duration during the sprint showed a mean absolute error of $4.3 + / - 3.2ms$ [41]. In terms of accuracy, these results outperform the results of related work that used lower sampling frequencies [189]. The mean absolute difference towards the reference system is consistent with that of related work [190].

In the *Long Sprint* analysis, athletes where continuously monitored during sprint on a 400m competition track. The temporal structure of the typical performance parameters, i.e., stance duration and step frequency, was measured to analyze the fatigue level of the athletes. The results showed that stance durations increased and step frequencies decreased steadily during the sprint and reached their extreme values in the last section of the track. This coherence shows how fatigue affects the performance of athletes, moreover as the results showed that the fatigue level is different for athletes with different training levels [40].

The *SpoSeNs Node* has furthermore used for pilot studies such as analyzing hurdle sprint and volleyball spikes as well as for optimizing gaze strategies during the flight of complex acrobatic maneuvers [46]. In the volleyball study the determined jumping heights showed a mean difference of about 1cm compared to the reference system, which is comparable to the results of related work [171].

Discussion and Conclusion

All scientific studies that have been conducted with the *SpoSeNs Node* have proven that this measurement method of employing a wearable sensor node delivers reliable results and enables powerful analysis. Furthermore, due to its versatile usage such that it is not bound to any environmental conditions, in-field usage is achievable with continuous monitoring and thereby measurement range and duration restrictions of related approaches [191, 192] can be overcome. Besides that, the ability of real-time feedback, enabled by embedded analysis algorithms and wireless communication, delivers just-in-time performance measurements of the athlete to the coach during training and competition [41,46]. Therefore, the presented wearable sensor platform represents a considerable contribution for professional sports, as its employment enhances the quality of training that is required in elite sports [193].

3.2.2 Wireless Synchronization

The previously discussed wearable sensor node platform is designed to wireless communicate with an end device, e.g., a smartphone, to report real-time performance parameters as a feedback to the current sports performance of the athlete. This approach is sufficient for sports disciplines, where a single sensor device is sufficient for the movement analysis. However, specific sports disciplines also require to measure the body movements at different locations of the athlete's body. Such distributed sensor data acquisitions require a precise time synchronization of the among the sensor units for a reliable performance analysis. Therefore, the next two sections present new synchronization approaches for wireless technologies.

Bluetooth Low Energy[4]

The new wearable sensor device *SpoSeNs Node*, which is discussed in the previous section, is equipped with the new BLE wireless communication standard. At the time of this research, the BLE standard has just been released as a subset of Bluetooth 4.0 specification.

BLE is specifically designed to ensure long battery lifetimes for battery-powered devices. Therefore, the base concepts guarantee a very low energy demand despite wireless communication. The main concept is that the radio peripheral can rest in idle mode most of the time and only is activated in specific time slots, called connection events. The time interval between these connection events is called connection interval t_{CI}. This interval is determined upon a connection establishment between the initiating device, referred to as central device, and the peripheral device. Thus, a precise timing schedule in the link layer is required, that must not be interfered by the user application. Therefore, data transfer requests by the user application are buffered and automatically processed upon the next connection event. As a result, the application program cannot determine the point in time when the radio becomes active. This ensures a reliable connection handling that is independent from the application. However, precise synchronization among connected devices become quite a challenge, which is why most approaches of related work showed poor synchronization accuracy as discussed in Section 3.1.2.

Based on the results of related work as discussed in Section 3.1.2, software-based approaches only lead to sufficient results when the BLE stack is customized to enable a direct radio peripheral access [73]. However, this results in a solution that is not compatible with the BLE standard. As a solution, this work presents a hardware-based approach that detects the radio activity by its characteristic current demand pattern. Since the timing of the link layer ensures that the radio peripherals of two connected devices are activated concurrently, this pattern occurs at both devices at the same point in time with a minimal deviation. To detect this pattern, the proposed hardware design considers a shunt resistor in the voltage

[4]This section and the related graphics have already partially been published in [39]

supply path. As shown on the left hand side in Figure 3.5, the voltage across the shunt is amplified and connected to a comparator inside the *System on Chip* (SoC).

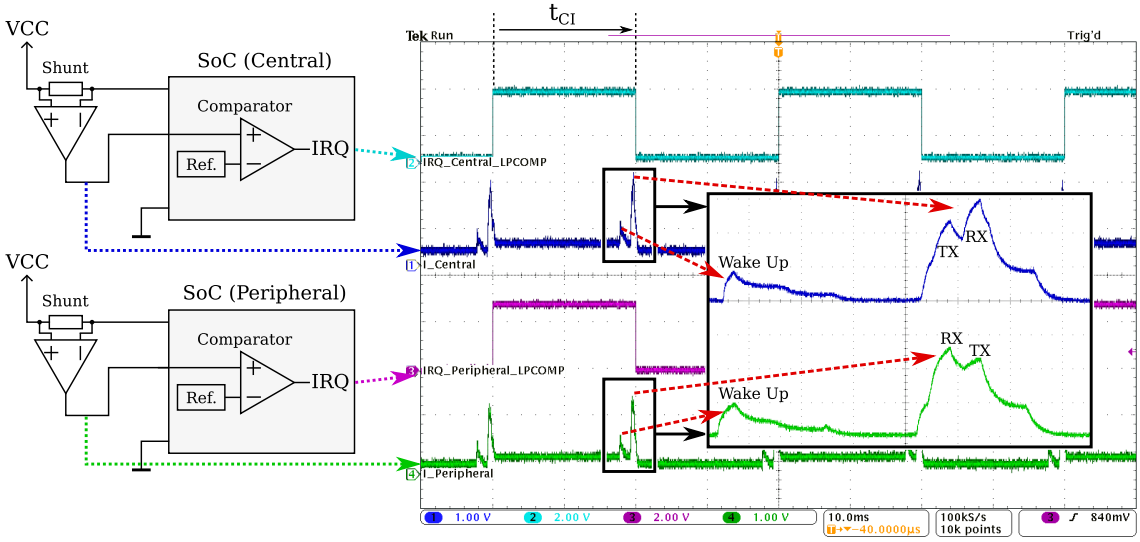

Figure 3.5: Proposed circuitry for detecting the current peaks, drawn by the BLE radio in the SoC, together with the comparator *interrupt request* (IRQ) outputs of two interconnected devices.

This way, an IRQ is fired by the comparator in both devices, displayed as signal edges in the cyan and magenta-colored signals of the oscilloscope graph in the right hand side of Figure 3.5. The enlarged area inside the oscilloscope graph zooms in the signal form of the amplified shunt voltage during a connection event. Here, the first peak is caused by the wake-up of the SoC from sleep mode. The next two peaks indicate the activation of the radio peripheral, where the larger peak represents the activation of the receive (RX) circuitry and the smaller one the activation of the transmit (TX) circuitry. Thus the second peaks indicate a transmission from the central to the peripheral device, and the last peaks a transmission in die opposite direction, which is the official transmission order defined in the Bluetooth specification and described in the appendix C.1.

The delay between the comparator IRQs in both devices have been measured and statistically analyzed. The outcome has been evaluated against the accuracy of the main clock source for the SoC and the chosen connection interval t_{CI}, as these potentially affect the timing of the BLE radio. Clock sources like the on-chip RC oscillator, which are less accurate than external crystal-based oscillators (XTAL), require an earlier radio activation at the peripheral device in order to compensate the possibly increased frequency error as shown in Figure 3.6. This graphic shows the time delays between the current demand curves of the interconnected central and the peripheral device, $i_{Central}$ and $i_{Peripheral}$, as represented in the respective amplified shunt voltages V_{Shunt} together with the reference voltage of the comparator V_{Ref}. Longer connection intervals can lead to an accumulation of the time base error, as the idle time of the device is longer. Therefore, it can be expected that the deviation of the IRQ occurrence in both devices becomes larger for longer connection intervals

and less accurate clock sources. The solution to these issues is to not use the rising edge of the current peak, which occurrence depends on the above-mentioned factors, but to assess the falling edge. At the falling edge, the radio peripherals are disabled after a connection event, which means that the compensation of the inaccuracy has already been performed. Figure 3.6 shows that the deviation of $t_{\Delta RE}$ is usually larger than the deviation of $t_{\Delta FE}$ and furthermore depends on the accuracy of the clock source.

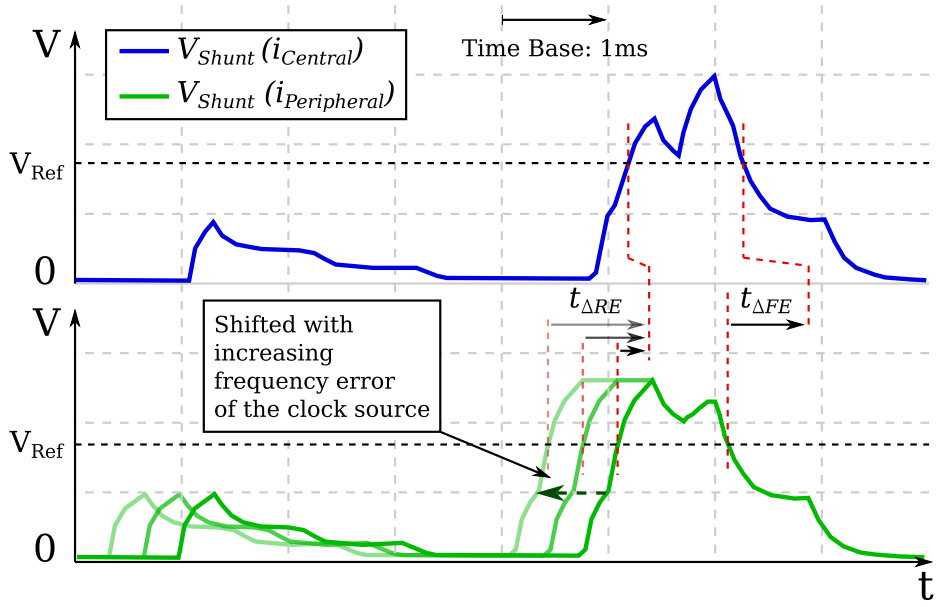

Figure 3.6: The time delays between the IRQs for rising ($t_{\Delta RE}$) and falling edge ($t_{\Delta FE}$) trigger.

Therefore, the time delay between both IRQs is defined as $t_{IRQ\ delay} = t_{\Delta FE}$ for the proposed method and its deviation determines the achievable synchronization accuracy. The frequency error ϵ of crystal-based oscillators is defined in ppm and specific to the employed crystal. The frequency error of the on-chip RC oscillator is defined in the data sheet of the SoC. In order to compensate this error, the designer can chose between different calibration intervals t_{cal} for the RC oscillator.

The following results in Figure 3.7 and Figure 3.8 show that when evaluating the falling edge of the current peak, the variation of the interrupt occurrence does neither depend on the connection interval nor on the accuracy of the clock source.

For the longer connection intervals, less data sets have been recorded than for the shorter intervals, therefore these t_{CI} values exhibit less occurrences than the others. The results show, that neither the clock source nor the connection interval affects the results when choosing the falling edge as IRQ trigger. Therefore, the standard deviation of $t_{IRQ\ delay}$ in all settings is the same and has been determined to $0.9\mu s$. This standard deviation determines the accuracy of wireless synchronization applications that exploit the presented methodology.

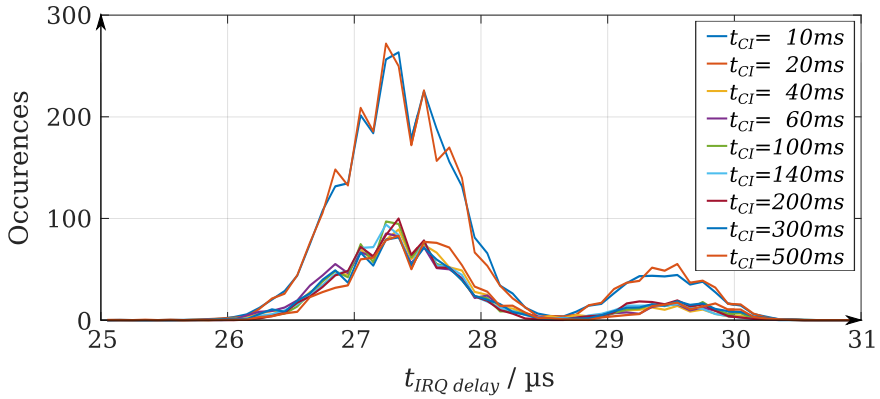

Figure 3.7: Synchronization accuracy with respect to different connection intervals t_{CI}.

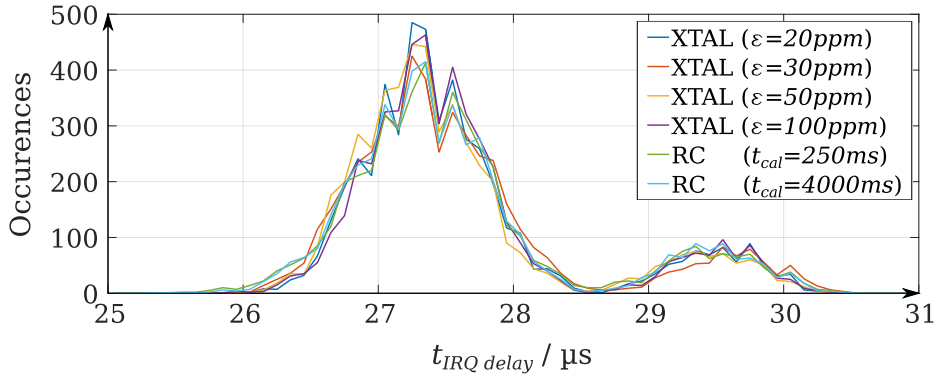

Figure 3.8: Synchronization accuracy with respect to different clock sources (crystal-based (XTAL) and RC oscillator-based (RC)).

The presented method leads to significantly lower synchronization errors than the approaches of related work, while being completely compliant with the BLE standard. Table 3.2 shows the result of the proposed principle compared to the approaches of related work.

Principle	Accuracy (std. deviation / average latency)
This work	$\pm 0.9 \mu s$
CheepSync [72], [73]	$10 \mu s$[a]
Bideaux et al. 1 [71]	$\pm 0.41 ms$
Bideaux et al. 2 [71]	$\pm 14.9 \mu s$
Ghoshdastider et al. [194]	$37.781 ms$
Ghoshdastider et al. [70]	$45.071 ms$

[a] Custom, nonstandard BLE stack

Table 3.2: Synchronization accuracy results of the presented methodology compared to related work.

Real-Time Body Area Network *SyRealNET*

In the previous section, a new wireless time synchronization method has been proposed for BLE-based *Personal Area Networks* (PANs) to overcome the protocol-intrinsic separation of physical and application layer that hampers wireless time synchronization applications. This method generates hardware time stamps that can be used for a time synchronization between two interconnected devices, i.e., a central and a peripheral device. In order to establish a synchronized network with multiple communication parties than two, the respective central device needs to hold multiple connections to multiple peripheral devices and needs to distinguish which hardware time stamp is related to which of its connected slaves. However, there is no reliable possibility in the central device to assign specific connection events, as they are reflected in the current peaks, to a specific peripheral device.

Therefore, a new approach is presented in this section to design a synchronized network with multiple communication parties. This approach is not based on the BLE protocol, instead a proprietary protocol for a BAN has been designed.

To evaluate this protocol, a dedicated hardware was designed that forms an extensive wireless BAN with up to ten wireless sensor devices. The BAN is designed for a low power operation to enable battery-powered devices and due to its wireless connectivity it provides a high degree in flexibility and therefore can compete against related wire-based sensor systems on the market.

Wireless Body Area Network Topology

The wireless BAN consists of up to ten battery-powered sensor devices and one gateway device as shown in Figure 3.9.

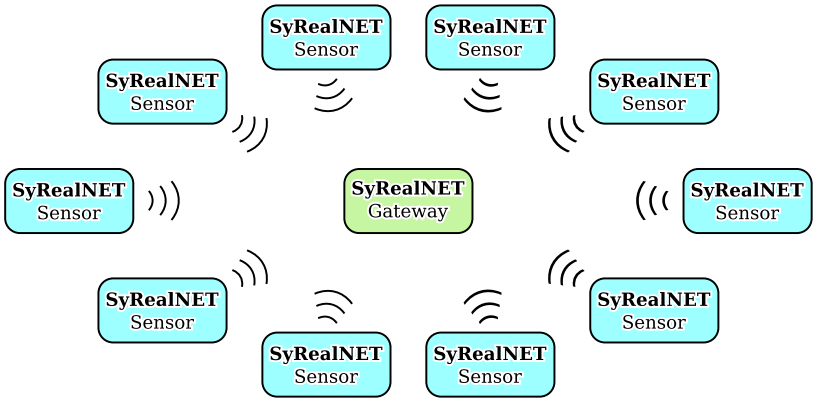

Figure 3.9: Star form topology of the proposed BAN with ten sensor devices and one gateway.

From a communication perspective, the gateway device acts as master and synchronization commissioner whereas the sensor devices act as slaves. The sensor devices receive a synchronization command all at once and trigger their internal sensing peripherals. Then, the sensors transmit their sensor values one after another to the gateway, who transfers the data sets to a host computer via USB.

Synchronization Principle

The employed synchronization principle is similar to the *IO-Link Wireless* protocol [195, 196]. *IO-Link Wireless* is an extension of the popular *IO-Link* protocol, which replaces the wired physical layer with a wireless data link. The wireless communication protocol is based on a master slave topology with *time division multiple access* (TDMA) access. The first step of each transfer cycle, the master transmits a multicast downlink frame to all slaves as shown in Figure 3.10. During the second step, the slaves transfer an uplink frame to the master one after another to the master. This way, an intrinsic time synchronization among all slaves is performed with the downlink by the master, similar to the *Reference Broadcast Synchronization* (RBS) [177] principle.

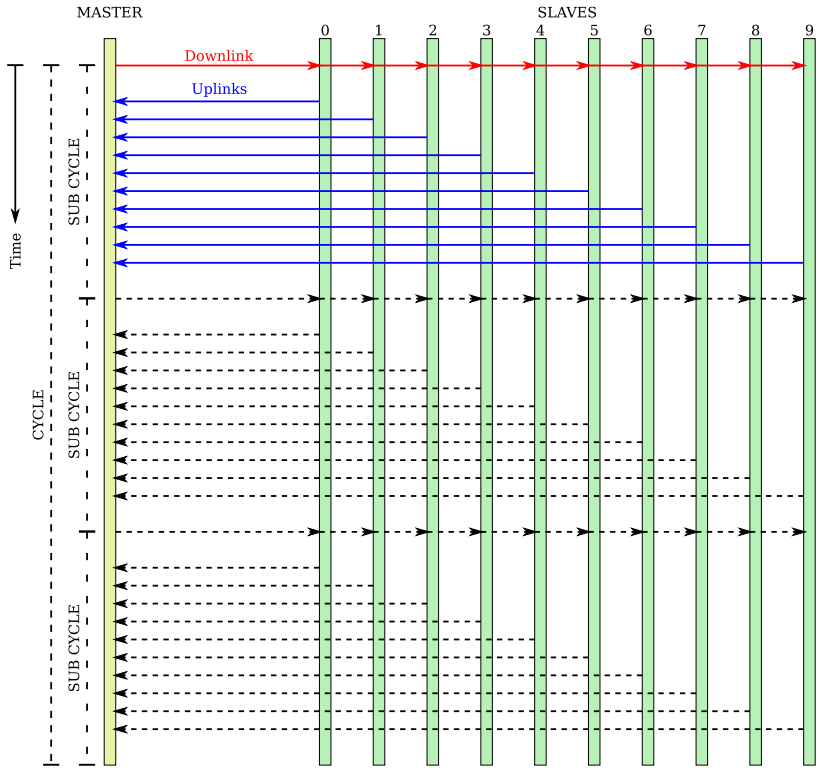

Figure 3.10: Each transfer cycle that starts with a downlink by the master followed by consecutive uplinks from all slave devices.

The synchronization error among all slaves in the *SyRealNET* is determined by the following error sources:

- Different wireless signal propagation times due to different distances from the master.
- Delays in the RX interrupt handling in the SoC.
- Wireless signal reception in the physical layer of the receiver chain.

The first error source is negligible due to the usually short distance differences between the slaves and the master in a BAN. The second error source has been minimized in the

embedded firmware by an interrupt driven design. As the interrupt handling is the same in each slave and is synchronous with the main clock frequency, the synchronization error of this is within the range of one clock cycle of the main oscillator, which runs at $48MHz$. The third error source is determined by the granularity of physical bit rate which is $2Mbps$, leading to a signal reception quantization error of $\frac{1}{2Mbps} = 500ns$. Since the third error source is orders of magnitudes higher than the other sources, the maximum synchronization error is determined by this error source.

Sensor Data Acquisition

Upon reception of the downlink frame by the master, each slave device triggers a measurement of its internal IMU as shown in Figure 3.11.

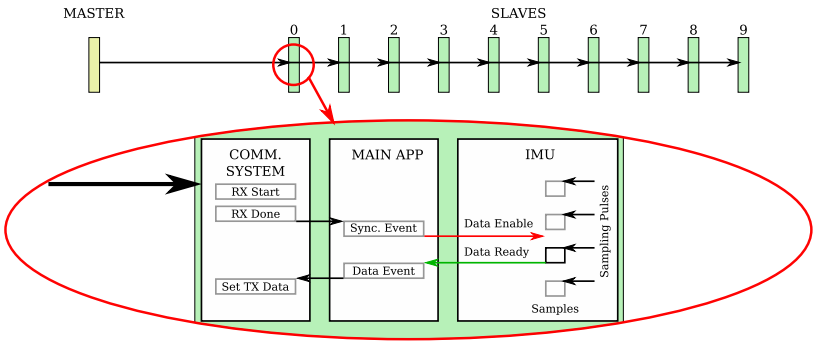

Figure 3.11: Sensor data acquisition of the IMU, synchronous with the network synchronization.

This way, the sample impulse for the sensor data acquisition takes place in each slave device at the same point in time. The deviation is determined by the configured sample rate of the IMU, which cannot be synchronized with the wireless communication system since it's internally controlled by the employed IMU. Therefore, the highest sample rate leads to the lowest deviation.

SyRealNET System Design

Figure 3.12 shows the hardware concept of the *SyRealNET* node that mainly consists of an ARM Cortex M4-based *radio frequency* (RF)-SoC, a rechargeable battery, and an IMU MEMS sensor. The vibration motor is intended for feedback mechanisms during rehabilitation therapy, which is however subject to current research.

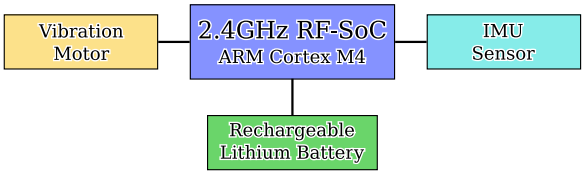

Figure 3.12: Hardware concept of the *SyRealNET* Sensor.

As multiple sensor devices are intended to be worn on the body, the requirements for the system design have been even stricter than the platform presented in Section 3.2.1. They

needed to be significantly more lightweight, such that the main contributor to the weight — the battery — had to be changed to a very small type. The chosen type only weights a few grams and comes with a capacity as low as $70mAh$. In return, this restricts the energy consumption of the sensor device and requires an ultra-low power operation. This has mainly been achieved by a dedicated embedded software architecture that exploits the SoC-specific programmable peripheral interconnection features for an event-and-task-driven event and action management. Figure 3.13 shows the architecture of this peripheral interconnection feature, which interconnects events of peripherals via programmable channels with tasks of other peripherals and thereby enable a control flow without *Central Processing Unit* (CPU) intervention.

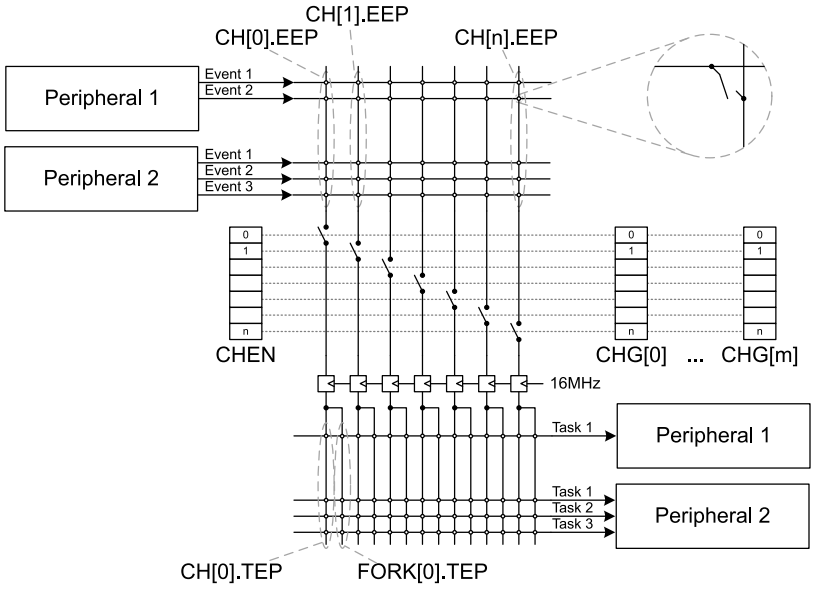

Figure 3.13: The basic architecture of the programmable peripheral interconnect (PPI) (source: Nordic Semiconductor).

Temporary idle phases of peripherals are avoided by powering them down once the scheduled idle time exceeds the respective wake-up time. To ensure a precise timing despite a frequent power-down and wake-up procedures, the wake-up times for each peripheral are considered in the event schedule. In general, the main context of the CPU is always off, which ensures an ultra-low power operation of the desired synchronization operation. More details about the implementation and the peripheral interconnections are documented in the *SyRealNET* documentation [197].

The system hardware has been designed as low cost system that only requires off-the-shelf electronic components. A proper enclosure as shown in Figure 3.14 enables a flexible attachment on the body with belts. Due to its low weight of $\approx 31g$, any negative influence on the mobility is minimized.

Figure 3.14: Enclosure of the sensor device for flexible attachment.

Discussion and Conclusion

The presented new synchronization methods enable a low power wireless synchronization of distributed WSNs. The first method is dedicated to synchronize two devices and relies on the BLE standard, whereas the second is intended for multiple devices and proposes a dedicated wireless synchronization protocol. Compared to approaches by related work, these methods exhibit an increased synchronization accuracy.

Chapter 4

Low Power High-End Security for the Internet of Things

> ”Privacy is the right to a free mind.”
>
> Edward Snowden

Like the previous chapter, this chapter discusses mobile embedded systems as indicated in Table 4.1. However, whereas the previous chapter specifically addressed portable and wearable devices that are powered by rechargeable batteries, this chapter focuses on energy-specific design challenges of devices that are powered by primary battery cells.

Autonomous Embedded Systems	**Mobile Embedded Systems**	Stationary Embedded Systems

Table 4.1: This chapter addresses mobile embedded systems. More specifically devices that are powered by primary battery cells.

Primary battery cells are the usual power supply for non-portable and fixed electronic devices, for which no local power supply, i.e., energy grid, is available. In contrast to portable and rechargeable devices in the previous chapter, non-portable devices require very long battery lifetimes in the range of years to ensure an adequate usability. This applies specifically to locally fixed and stationary-mounted *Internet of Things* (IoT) edge devices, where every battery replacement represents an additional organizational and economical effort that must be avoided. As IoT edge devices are connected to the internet, a further critical aspect is the security of the communication where high-end security mechanisms are inevitable to be resistant to attacks. This need for — typically computationally and energy-intensive — security algorithms leads to the most critical microelectronic design challenges, since it is contrary to the requirement for long battery lifetime, to which a low energy demand is a prerequisite. This becomes even more critical due to the fact that embedded computing

platforms of IoT edge devices typically come with limited computing performance compared to conventional internet-connected devices like desktop *Personal Computers* (PCs) and servers.

Key aspects to a successful IoT are flexible and affordable communication infrastructures as well as low development costs for the *things*, i.e., IoT edge devices. Therefore, as far as the edge devices are concerned, the success of the IoT depends on the ability of employing low cost edge devices that can operate from small batteries for a long time. This leads to a conflict for IoT infrastructures that simultaneously aim at a high level of security. On the one hand, the edge devices come with several restrictions due to their small energy budget and limited computing performance compared to a personal computer due to their embedded *Microcontroller Units* (MCUs). On the other hand, high-level security protocols like *Transport Layer Security* (TLS) come with computationally intensive algorithms, which similarly applies to the protocol handling for network protocols like *Internet Protocol* (IP) and transport protocols like *Transmission Control Protocol* (TCP).

This chapter is structured as follows: Section 4.1 covers the technical background and challenges regarding IoT connectivity, general IP-based internet security and the state-of-the-art in security approaches for resource-constrained IoT edge devices. Section 4.2 presents the new scientific contributions for ultra-low power system design of IP-enabled IoT edge devices and shows how this allows for an end-to-end TLS-encrypted communication.

4.1 Technical Background, Scientific Challenges and Related Work

The concept of IoT has a broad variety of technology topics and aspects: Data collection, processing and analysis, cloud connectivity, system design of edge devices, gateways and servers, etc. This work is dedicated to the cloud connectivity, i.e., how edge devices are physically connected to the internet, and the system design of secure and low power edge devices.

4.1.1 IoT Connectivity

As far as the cloud connectivity is concerned, a large variety of solutions has been developed during the last years. Figure 4.1 shows some popular connectivity topologies that have been pushed by companies and standardization initiatives:

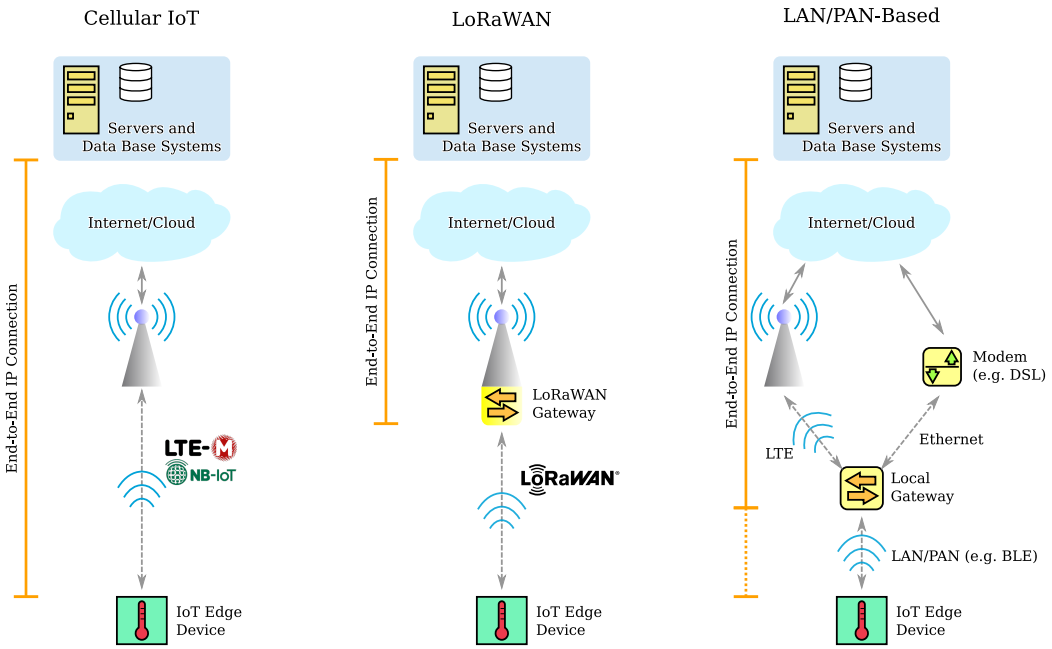

Figure 4.1: Variants of Low Power IoT connectivity for battery-powered edge devices (technology logo sources: [198, 199]).

Each topology has its specific strengths and weaknesses in terms of flexibility of usage, energy requirements for the edge devices, implementation complexity and security features.

The basic idea behind *Cellular IoT* (CIoT) is to connect everything directly via *Wide Area Networks* (WANs) to cellular base stations and therewith to the internet without needing a smartphone or a gateway. Therefore, cellular communication standards have been adopted to address the limited energy budget of constrained edge devices. *Long Term Evolution* (LTE), is one of the most popular standard besides *5G* that is released by the *3rd*

Generation Partnership Project (3GPP) [198], a global initiative that develops and manages cellular standards. Depending on its performance, LTE is split up in categories that differ in bandwidth, *Multiple Input Multiple Output* (MIMO) capabilities, modulation, etc. Originally, LTE has been developed to speed up the cellular internet and cloud connectivity, without specifically addressing low power operation for constrained edge devices. To enable low power with the cellular LTE infrastructure, dedicated LTE categories have been released by the 3GPP that differ to conventional categories in less bandwidth, lower-level modulation, half-duplex transfer and omitted MIMO support. Exemplary categories dedicated to low power CIoT are *Cat-M1*, also known as *LTE-MTC (Machine Type Communication)* (LTE-M), and *Cat-NB1*, also known as *Narrowband-IoT* (NB-IoT). Compared to NB-IoT, LTE-M comes with a higher bandwidth and throughput as well as a lower latency. In return, NB-IoT offers a higher range to the base station and provides a better signal penetration in buildings.

CIoT provides a flexible system usage and deployment, however it requires cellular coverage that highly depends on the region and the national service providers. Furthermore, every communication party, i.e., *thing*, must be equipped with a *Subscriber Identity Module* (SIM) for the required network authentication with the base station, which results in individual service fees per device.

Alternative WAN-based topologies use proprietary network infrastructures (e.g., SIGFOX [200]) or employ dedicated gateways to route the traffic to the internet (e.g., *Long Range WAN* (LoRa-WAN)).

Transverse approaches to WAN-based infrastructures employ shorter range networks like *Local Area Networks* (LANs) or *Personal Area Networks* (PANs) to sum up multiple edge device connections to one single internet endpoint by using a local gateway device. This way, network operation fees can be reduced significantly or even existing wired internet endpoints can be used. Furthermore, the system design of the edge devices can be less complex, as only local network connectivity is required. Shorter ranges of wireless transfer protocols also come with lower energy demands in relation to the bandwidth compared with WANs. However, the requirement of local gateway devices results in additional device development and maintenance, which can lead to serious issues especially regarding security aspects. Each security-relevant device poses a potential security risk on the IoT system.

4.1.2 Security

The seriousness of security-related issues and the vulnerability of current IoT infrastructures became obvious due to many data leaks and reports of bot net misuse of IoT devices that became public in the recent years. Therefore, security has been valued as one of the biggest challenges in order for the IoT concept to be successful [201]. The relevance for high-end security ranges throughout the entire field of applications, ranging from privacy

issues in consumer-oriented IoT systems to industrial applications (*Industrial Internet of Things* (IIoT)). IIoT concepts, e.g., enable industrial machinery to be interconnected with cloud services to optimize industrial processes. Especially there, transferred data is often related to critical and safety-relevant infrastructures.

The state-of-the-art data transport security protocol for the internet is TLS. Prior to a new data connection between two communication parties, a handshake procedure is conducted to establish a secure connection. This handshake procedure is built on mutual authentication of the communication parties and secure session key exchange procedure. The authentication process prevents *Man-in-the-Middle* attacks by using cryptographic methods, such as *Digital Signature Algorithm* (DSA), to ensure that the identity of every communication party is unique and cannot be falsified or stolen. The key exchange procedure utilizes cryptographic methods like *Diffie-Hellman* (DH) to securely transfer an encryption key over an insecure channel. An additional level of security can be achieved by establishing a new encryption key upon every connection. This level is called *Perfect Forward Secrecy* (PFS) and it employs the DH protocol with ephemeral keys (*Diffie-Hellman ephemeral* (DHE)). Figure 4.2 shows the simplified handshake procedure.

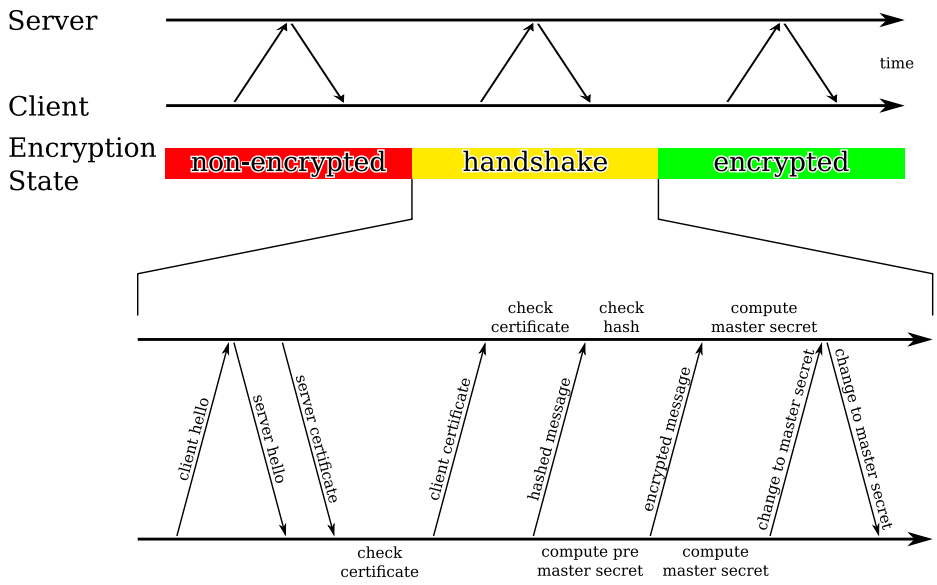

Figure 4.2: TLS handshake as initial procedure for an encrypted communication.

After the handshake procedure, all further communication is encrypted with the session key. The authentication is based on certificates and relies on an existing global *Public Key Infrastructure* (PKI) with commonly trusted certificate authorities as root of trust [202]. This PKI is existing in the internet for many decades and provides trust and security to all well-known security-relevant services like online banking. TLS relies on the TCP to establish a secure data transfer via an end-to-end IP-based connection. The achievable data transfer security levels of TLS are valued as very high and certified by various institutions. Therefore, it is broadly applied in all common internet communications for many decades.

4.1.3 IoT Security for Constrained Edge Devices

The conventional cryptographic algorithms employed for authentication and key exchange are based on *Rivest–Shamir–Adleman* (RSA) cryptography, which works with key lengths of at least 2048 Bit. These key lengths place high requirements on communication throughput, computing performance and local memory. In order to reduce the related overheads for the embedded computer platform of the edge device, *Elliptic Curve Cryptography* (ECC) has emerged as an alternative to RSA. ECC reduces the computational effort significantly with respect to classical RSA algorithms [203], especially regarding the key exchange [204] as the key sizes are reduced significantly by keeping the same level of security.

However, despite the reduced computational effort of ECC, TLS is usually not employed in IoT systems to secure data transfer between low power edge devices and servers. First because the computational requirements are still too high for the limited computing capacity of low power edge devices [74], and, second, due to missing end-to-end communication between edge device and server. The reason for the latter issue is mostly because of the required gateways that bridge the edge device communication to the internet, but also because internet-compliant protocols like IP are considered as too bulky and too computationally intensive for resource-constrained embedded platforms [74–77]. The most likely TLS connection only exists between the gateway and the server like, e.g., in LoRa-WAN and other proprietary solutions.

The importance of advanced security standards for the IoT is a broadly accepted fact and has been intensively studied in the last years [76, 201, 205]. Many proposals have been published in the recent years to manage and ensure trustworthiness in the heterogeneous IoT world [206]. However, it seems to remain an unsolved challenge, especially due to the vast heterogeneity of existing IoT devices [207]. The fact that most deployed IoT devices are not as simple configurable and upgradable as conventional IT systems emphasizes the requirements for these devices to provide security by design [74].

As a solution, related work came up with lightweight alternatives to TLS like energy-efficient Datagram TLS (eeDTLS) [78], E-Lithe [79], Compact TLS (CTLS) [208] and many others [80, 81, 209, 210]. Common disadvantage of lightweight alternatives is that they do not offer the identical level of trust as TLS. Furthermore, their usage leads to compatibility issues, as these lightweight protocols are mostly not supported by public internet services and major cloud system vendors [211–213].

In order to establish a highly secure IoT infrastructure that allows for using low power and low complex edge devices, a PAN/LAN network technology must be employed that can efficiently transfer IP data transfer. The popular *Wireless LAN* (WLAN) or *wireless fidelity (WLAN)* (WiFi) technology however is known for extensive energy demand and thus is not suited for battery-powered edge devices. Regarding PAN technologies, the *IPv6 over Low Power Wireless Personal Area Network* (6LoWPAN) [214] standard has been established to enable an *Internet Protocol v6* (IPv6)-based data transfer over low power and bandwidth constrained communication protocols. 6LoWPAN introduces several concepts

for efficient IP data transfer, such as packet compression and fragmentation. Established implementations are provided by *Bluetooth Low Energy* (BLE) and IEEE 802.15.4 [215]-based protocols like *Thread* and *ZigBee*, whereas BLE has been shown to be more robust to obstacles [216] and to outperform IEEE 802.15.4 in terms of energy efficiency in general [217, 218] and specifically for 6LoWPAN [219].

However, at the time of this research, related work had not yet presented a design approach for a battery-powered IoT edge device that offers end-to-end IPv6 connectivity together with PFS-grade TLS encryption at reasonable battery lifetimes.

4.2 New Contributions

This section presents the new scientific contributions of this thesis in the research field of mobile embedded systems, powered by primary battery cells. It specifically addresses the challenges to enable sophisticated security features on non-portable, battery-powered IoT edge devices. The contributions are divides into two topics:

- Section 4.2.1 presents a new platform design for an ultra-low power IoT edge device. To ensure robustness against common issues in third-party software components, e.g., IP stack, a dedicated hardware-based mechanism for a stateless operation is designed.
- Section 4.2.2 addresses the design challenges between ensuring long battery lifetimes and providing high-end security at the same time. As sophisticated security protocols like TLS rely on computationally intensive security algorithms, they are known to be energy-intensive and therefore affect the battery lifetime. This work shows how an end-to-end TLS encryption becomes feasible for battery-powered devices that guarantees years of battery lifetime under certain assumptions and constrains. Therefore, a specific IoT infrastructure proposed as a new indicatory approach towards a secure and ultra-low power IoT.

4.2.1 Ultra-Low Power IoT Edge Device Platform Design

The hardware platform for the ultra-low power IoT edge device is constructed around the *nRF52840*, an ARM Cortex M4-based *System on Chip* (SoC) by Nordic Semiconductor [220] with an integrated BLE-compliant 2.4 GHz radio. This SoC comes with a *Software Development Kit* (SDK) and BLE software stack that supports 6LoWPAN standard to transport IPv6 via BLE. The advanced event and interrupt management system as well as the different power modes of this SoC allow for a low duty cycle-operation. Therefore, the node is in power-down mode most of the time and only activated in predefined time intervals or on specific events. This reduces the overall energy demand significantly in contrast to conventional computing with massive *Central Processing Unit* (CPU) intervention. On wake-up, the next active cycle starts, in which sensor data is obtained and transferred to the server. After successful data transfer, power-down mode is entered again.

Further components of the hardware platform are different types of sensors as shown in Figure 4.3.

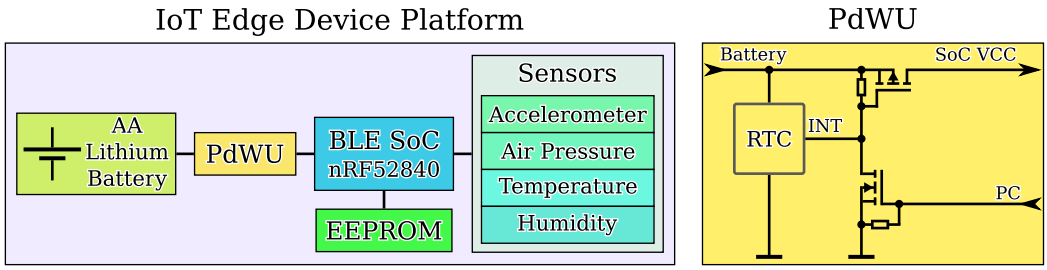

Figure 4.3: Hardware platform of the IoT edge device with a hardware-based stateless operation mechanism.

So far, this hardware design follows a standard approach for IoT edge devices used for environmental sensing applications. However, early evaluations of the system revealed the necessity of additional functionality, as the embedded software was crashing after some days of operation. Some reasons for this have been located in employed third-party software modules such as the SDK, IP stack or TLS implementation, of which many were in an experimental stage of development at the time of this research. Other reasons arose from timeout issues in the server software, which is usually not optimized to communicate with resource-constrained edge devices and the respectively longer packet processing delays. In order to provide a high level of robustness against such unforeseen software failures, the platform design was extended with a hardware-based mechanism that enables a stateless operation.

The stateless operation mechanism was realized by a custom *power-down and wake-up unit* (PdWU), which introduces a power gating for the main SoC. This PdWU involves a self-killing circuitry for the SoC and outsources the wake-up control to an external control instance. The external control instance is realized by a *Real-Time Clock* (RTC) chip that is configured just before power-down is issued and wakes up the SoC by an interrupt.

This way, each new active cycle is started by a power-on reset, due to which the state of the embedded software is lost in the event of a malfunction and cannot affect the next active cycle. To log errors through the reset cycle, an *electrically erasable programmable read-only memory* (EEPROM) has been considered in the hardware platform as shown in Figure 4.3.

Via the power control (PC) line of the PdWU circuitry, the SoC keeps its power supply enabled during the active cycle and activates power gating for power-down (self-killing circuit). After a predefined period of time, the RTC activates the power supply for the SoC again by firing a low level interrupt. In power-down, the only continuously powered device is the RTC, which results in power-down currents of only a couple hundred nanoamperes.

The hardware comes as a small *printed circuit board* (PCB) as shown in Figure 4.4. It is powered by a single lithium AA cell as energy source, leading to small dimension of the edge device that ensures a flexible deployment and broad application range and comes at reasonable costs, as only off-the-shelf components have been used for the hardware design.

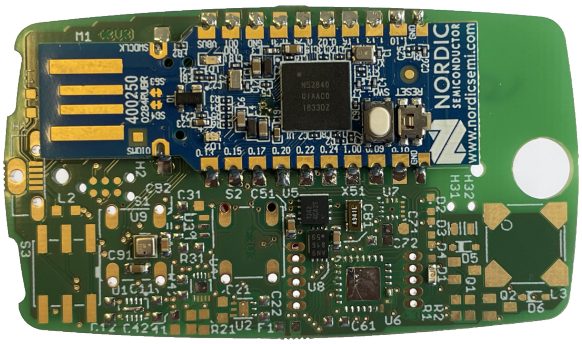

Figure 4.4: Hardware of the BLE-based IoT edge device, equipped with environmental sensors and further capabilities for additional components.

The following section presents the evaluation results of the TLS-security features, which was implemented for the IoT edge device platform described in this section.

4.2.2 End-to-End TLS with Perfect Forward Secrecy

Proposed IoT Infrastructure

The basic concept of TLS security is an end-to-end connection between the communication parties without intermediate stations that are involved in the handshake process. Therefore, a true end-to-end IP connection can be seen as a prerequisite for a truly secure TLS-based encrypted communication. As discussed before, many PAN protocols support IP-based transfer by, e.g., employing the 6LoWPAN standard. However, most of these protocols do not support TCP transfer but only *User Datagram Protocol* (UDP), due to restrictions in protocol handling and computational complexity. In contrast, since version 4.2, BLE supports full IP transfer with 6LoWPAN [221], supporting both TCP and UDP. Therefore, in this work, BLE was employed as low power PAN to interconnect the IoT edge devices and enable them as IP and TCP communication party. The proposed IoT infrastructure is based on the IoT communication concept by Nordic Semiconductor [220], a well-known manufacturer of BLE SoCs. Figure 4.5 shows the resulting communication topology of the proposed BLE-based IoT infrastructure.

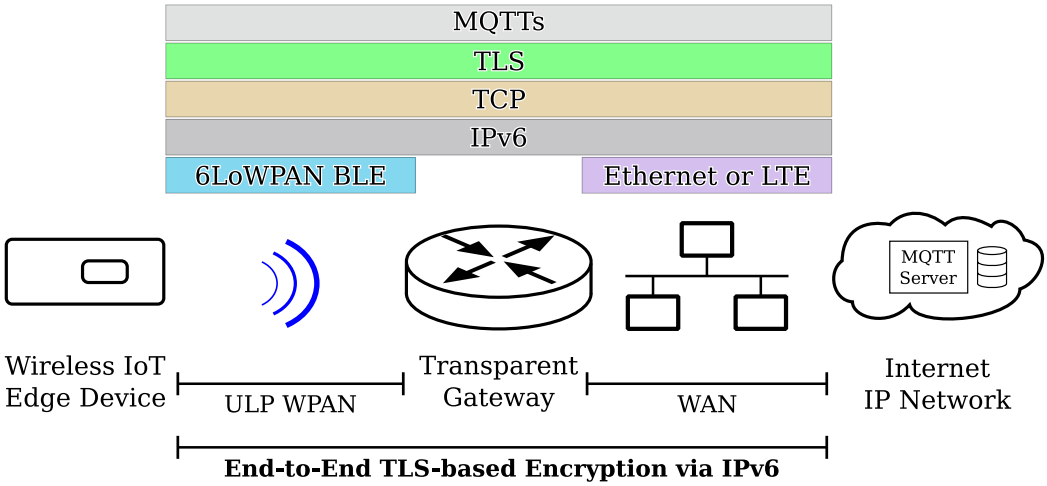

Figure 4.5: Proposed BLE-based IoT infrastructure.

As application protocol, the lightweight *Message Queuing Telemetry Transport* (MQTT) protocol was employed, as it is valued as the most promising TCP-based application protocol for the IoT [222, 223]. MQTT therefore became the de facto standard interface for most commercial cloud vendors for IoT solutions. On the server side, a *mosquitto* [224] was used as MQTT broker server.

The gateway is based on the popular small Linux computer Raspberry Pi 3 [225] and is configured with an optimized setup described by Nordic Semiconductor [226]. It is implemented as a transparent gateway, which only routes the IP transfer between the different physical layers of the PAN and theWAN. This way, the gateway does not touch any layer above the IP layer, thereby enabling a true end-to-end IP connection between edge device

and server. In return, the entire IP and upper layer protocol handling must be carried out in the edge device. Thus, the main design challenges for the proposed IoT infrastructure constitute the implementation of the bulky and computationally intensive protocol and security operations on the computationally constrained edge device.

The embedded software of the edge device is based on an implementation example for TLS by Nordic Semiconductor [227]. This basic example has been enhanced to provide a PFS-grade security level, i.e., full handshake functionality for mutual authentication and secure key exchange. The following third-party software modules have been employed:

- 6LoWPAN implementation for BLE: Nordic IoT-SDK
- IP stack for IPv6 connectivity: *lwIP* [228]
- TLS v1.2 protocol implementation: *mbedTLS* [229]
- Cryptographic library for the ECC-related computations: *Micro-ecc* [230]

The major design challenge was to efficiently map the resource-intensive modules like the IP stack and the TLS library to the limited resources of the SoC.

Evaluation and Results

The evaluation was focused on investigating the security-related overheads in terms of energy and latency. Therefore, the IoT edge device has been designed as a simple temperature logger that transfers only 10 bytes of user data to the server in a connection cycle. Potential use cases have been limited to environmental sensing applications, in which no fast changes occur that have to be reported and where only a few data transfers per day are sufficient.

Figure 4.6 shows a typical profile of the current demand of the IoT edge device during a full connection cycle. A connection cycle start with establishing a BLE connection to the gateway, followed by the TLS handshake procedure and subsequent MQTT connection to the server that includes the transfer of the user data, and ends with the BLE disconnection from the gateway.

The individual cryptographic operations during the handshake are marked in Figure 4.6 as *ECDSA* for the mutual authentication and *ECDHE* for the key exchange. The current draw i_{RMS} exhibits a relatively high static current during these operations, which indicates that the CPU is running continuously for these several tens of seconds. This indicates that the cryptographic operations are the major contributor to the overall energy demand and computation latency. In the other regions, the current basically consists of multiple subsequent peaks, which are mainly caused by the BLE radio activity. The low static load indicates that the CPU load is relatively low in these regions and that the embedded software mainly operates interrupt and event-driven there.

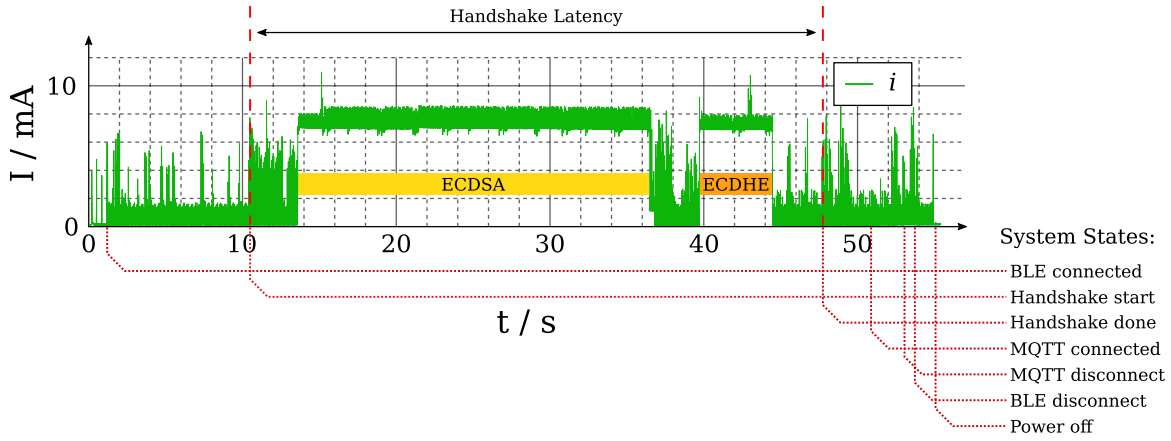

Figure 4.6: Characteristic current profile of a handshake and a consecutive MQTT data transport of the IoT edge device.

The measurements revealed that the average energy demand per handshake is be around $750mJ$ at a supply voltage of $3.6V$, and the handshake latency around $37s$. Within this latency, the ECC-related computations constitute about 80%, of which three fourth stem from DSA computations and one fourth from DHE. In contrast to conventional encrypted internet connections that reuse exchanged keys in further connections to save recurring handshake, the principle of PFS provides an improved level of security by performing a full handshake procedure upon each new connection. As this way once exchanged keys become invalid after each connection, PFS represents the most promising protection for IoT systems from their edge devices being compromised by attackers. But the reported results make clear that this comes at the cost of a high energy and latency overhead. Therefore, the achievable battery lifetime of the proposed IoT edge device mainly depends on the number of secure connection establishments per day as shown in Figure 4.7.

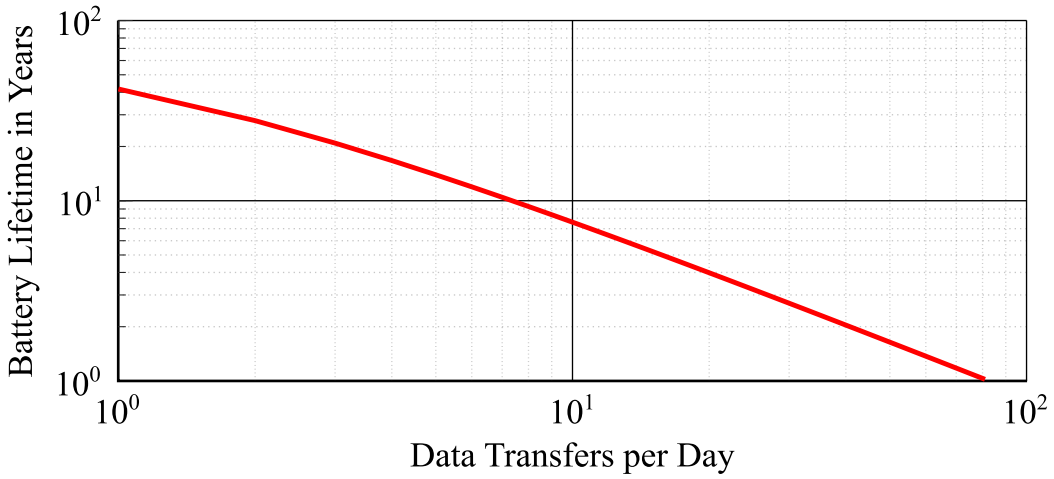

Figure 4.7: Battery lifetime over the number of TLS-encrypted MQTT connections per day of the proposed IoT edge device.

Further assumptions for this battery lifetime calculation are:

- Powered by a standard $3.6V$ Lithium Thionyl Chloride battery with 70% of its capacity of $2600mAh$ considered as usable.
- Power-down current of $1\mu A$ (worst case assumption for active PdWU).
- 10 bytes of user data transfer for each MQTT connection.

Discussion and Conclusion

To the best of our knowledge, this was the first work published that showed the feasibility of PFS-grade end-to-end TLS on resource-constrained, battery-powered IoT edge devices while guaranteeing years of battery lifetime from a single Lithium AA cell [52]. Therewith making the common belief of the necessity of lightweight alternatives to TLS for securing resource-constrained IoT edge devices obsolete.

Chapter 5

Embedded Model Predictive Control for Technical Infrastructure Facilities

$$\nexists\, a, b, c, n \in \mathbb{N}_{>0},\ n > 2 : a^n + b^n = c^n$$

Pierre de Fermat

The previous two chapters discussed mobile embedded systems, i.e., devices that are powered either by rechargeable or primary battery cells. This chapter, in contrast, addresses energy-specific design challenges of stationary embedded systems that are powered by the energy grid as indicated in Table 5.1.

Autonomous Embedded Systems	Mobile Embedded Systems	**Stationary Embedded Systems**

Table 5.1: This chapter addresses stationary embedded systems that are powered by the energy grid.

The energy grid, if locally available, is the most practical source of power for stationary electronic systems. Local availability is typically given in industrial environments, where — characteristic for the *Industry 4.0* concept — stationary embedded systems are employed to undertake key tasks in smart industrial environments like the control of actuators and processing plants. Due to its control stability and capability of handling many system components and parameters, *Model Predictive Control* (MPC) emerged as a promising control principle for industrial and technical infrastructure facilities. However, due to its computational complexity, the design of MPC controllers hold serious challenges, especially regarding energy-efficient implementations on configurable hardware platforms like *Field Programmable Gate Arrays* (FPGAs).

In this work, an MPC controller for a smart water infrastructure is designed, which comes with the specific requirement of a long prediction horizon[1]. The key challenge that

[1] The prediction horizon describes how far the control algorithm predicts the system behavior in the future.

is addressed is the implementation of the complex mathematical optimization algorithm of MPC in a parallel manner under the constraint of a low-cost FPGA with a comparatively low hardware resource count.

This chapter is structured as follows: Section 5.1 discusses the technical background and design challenges regarding MPC and FPGAs and presents embedded MPC implementation approaches for FPGAs as proposed by related work. Section 5.2 presents the new scientific contributions to significantly reduce the memory demand of MPCs controllers, which enables an embedded implementation of a long prediction horizon controller on a low power FPGA platform.

5.1 Technical Background, Scientific Challenges and Related Work

Modern infrastructure systems consist of various distributed entities, forming a grid of multiple interconnected units. For sustainable power supply, e.g., the idea of the smart grid emerged as one of the most promising concepts. Energy sources like power plants and wind turbines are connected for communication with energy distribution and storage facilities. In water distribution systems, pump stations are interconnected with water pipelines and valves, linking together sources like wells, water tanks and sinks. The coordination and control of these infrastructure systems is a challenging task, especially because a reliable operation is crucial. These control systems have to handle a large number of system components by following strict constraints, which points to the concept of *Model Predictive Control* (MPC) as appropriate control method.[2]

5.1.1 Model Predictive Control

MPC is based on a continuous mathematical optimization of the desired control objective by predicting the system behavior in the future. In comparison, conventional controllers control the system output to the desired value with a continuous feedback of the real output and try to minimize the error. Hence, usually a plain reactive principle is applied that is vulnerable to overshoots and short-term instabilities. In contrast, MPC incorporates a mathematical model of the system that is to be controlled. Using this model, the resulting system output can be calculated and predicted based on the input and the current system state. This allows to formulate the optimal input as an optimization problem (also referred to in the appendix D.1). As a result, the solution of the optimization problem leads to the best possible system behavior that compensates long term effects and strictly sticks to state constraints.

On the contrary, compared to conventional control systems, MPC usually incorporates large computational overheads due to computational complexity of the numerical optimization process and comes with relatively large memory demands. Therefore, for a long time MPC has been limited to slow applications like specific industrial processes and has required energy-intensive high-performance computers for execution. However, in the recent years, new achievements in computer hardware and the massive the deployment of *Graphics Processing Units* (GPUs) and FPGAs overcame the limited computing capacity of general purpose *Central Processing Units* (CPUs) and enabled MPC implementations on embedded computing platforms [85, 231]. Such implementations are more energy-efficient due to the massive parallelism of these platforms and the directly accessible high-bandwidth memory, which is essential for data-intensive tasks like MPC. Specifically highly customizable FPGAs platforms became more and more attractive for MPC due to

[2] Paragraph has already been published in [48]

their massive parallelism capabilities [232] that allow fast computations with a high energy efficiency [233–235].

5.1.2 Field Programmable Gate Arrays

FPGAs are programmable logic devices that consist of multiple *Configurable Logic Blocks* (CLBs) that contain basic units for memory, arithmetic and logic operations. These units can be configured and interconnected to form specific and complex hardware designs as combinatorical or sequential circuits. The available number of these units on an FPGA device are also known as hardware resources. Figure 5.1 shows the basic structure with CLBs as well as additional memory and computation blocks.

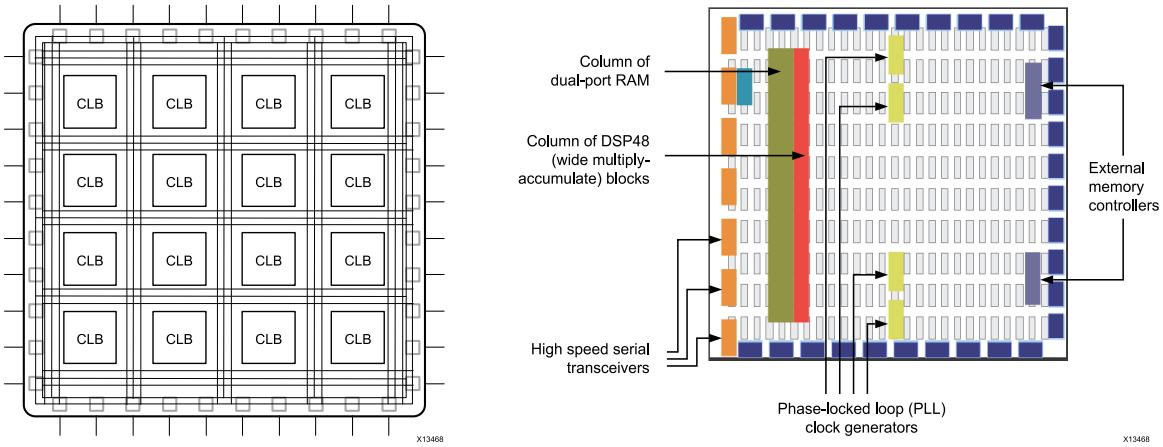

Figure 5.1: Internal structure of an FPGA: Basic architecture with distributed logic blocks (left) and additional memory and computation blocks (right) (source: Xilinx [236]).

In general, the following types of basic units are present on FPGA devices:

- *look-up tables* (LUTs): Their structure allows to, e.g., realize any logic function by a specific setting of the inputs of the multiplexer.
- *Flip Flops* (FFs): They are basic and fast memory cells that can be used as registers where, e.g., data is buffered between two clock cycles.
- *Block RAMs* (BRAMs): These are memory units that, e.g., can be used as random access memory blocks in the hardware design.
- *Digital Signal Processors* (DSPs) block: These units are *Arithmetic Logic Units* (ALUs). They are capable of performing arithmetic operations such as addition, multiplication and accumulation.

Figure 5.2 shows the internal structure of the LUTs, FFs and DSPs.

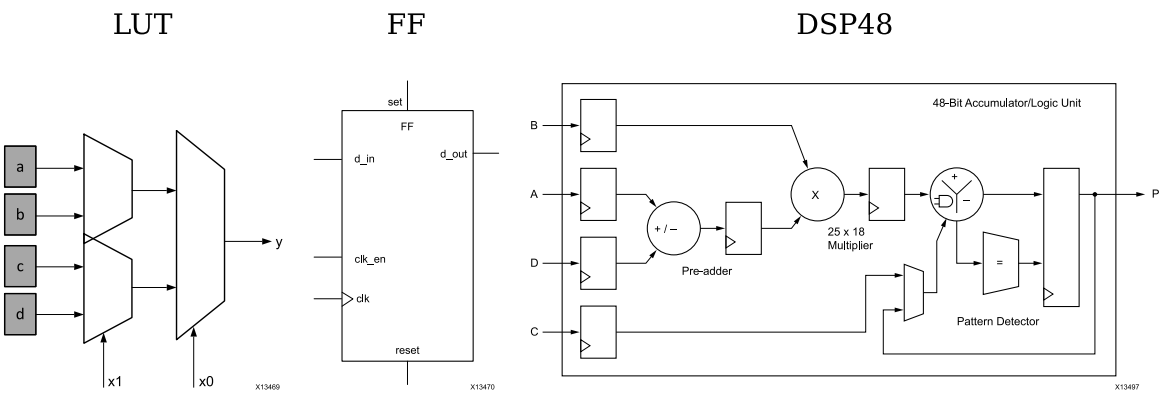

Figure 5.2: Internal structure of the units (source: Xilinx [236]).

In contrast to conventional CPU-based system design, hardware designers have a very high flexibility in designing the circuits for FPGAs. Whereas CPUs programs follow a sequential execution of instructions, hardware designs can make use of massive parallelization possibilities. Furthermore, CPUs are limited to a list of supported number formats, such as integer and floating point. In contrast, an FPGA allows to specify any number format presentation, even customized formats, which enables an optimal usage of existing hardware resources according to the given system requirements (also refer to the appendix D.3).

The amount of hardware resources, e.g., LUTs, FFs and memory, on FPGA devices is limited. Thus, the complexity of the system that can be implemented is also limited. Therefore, the size of MPC systems on FPGAs is restricted in most published applications. Primary only small to medium-scale MPC systems with relatively short prediction horizons have been explored.[3]

5.1.3 Embedded MPC on FPGAs

Control systems are usually modeled with the help of comprehensive math tools that allow straight forward controller design due to supportive libraries and dedicated environments. Matlab is such a tool that, e.g., enables sophisticated functional simulation and verification with provided toolboxes for control systems.

Hardware designs for FPGAs are typically implemented using a *Hardware Description Language* (HDL) like VHDL or Verilog. As HDLs usually operates at gate level, their usage requires a comprehensive knowledge of the underlying hardware architecture and typically induces a time consuming development process. In order to enable easier and faster design processes, alternative *High Level Synthesis* (HLS) methods have been released. HLS allows to use conventional programming languages for the hardware development like C/C++ or SystemC, instead of HDLs.

[3]Paragraph has already been published in [48]

Another approach is the use of code generators that generate HDL code out of other design languages. For instance, Mathworks [175] provides the Matlab *HDL Coder* to generate hardware designs out of Matlab source code.

However, generated HDL code tends to poor design quality regarding hardware resource utilization and energy efficiency. Matlab *HDL Coder* has been shown to synthesize resource-hungry hardware designs compared to Xilinx *Vivado HLS* [237], [238]. As a compromise, combining code generation with HLS has been proposed to generate resource-efficient implementations on resource-constrained FPGA systems [239].[4]

Related work presented implementations of embedded MPC for parallel computing platforms like FPGAs and GPUs [240–243]. These works differ in applied optimization algorithms, system applications and benchmark, performance requirements and system size. But, in general, all related MPC implementations profit from the massive parallel computing capabilities of these platforms. Regarding the achievable energy efficiency, FPGA-based implementations outperform implementations on both conventional computation platforms [234] and GPUs [244, 245], since they profit from a high degree of customization. Furthermore, FPGAs show better deterministic behavior than GPUs: Designs on GPUs can exhibit scheduling issues [246], whereas FPGA-based MPC designs can be employed for real-time applications [82–84]. The flexibility of FPGAs regarding customized number format representation has been exploited to enhance the performance of MPC designs by moving from standard floating-point to fixed-point formats [82, 83, 237, 247, 248]. As fixed-point arithmetic requires less complex computational units than floating-point arithmetic, the designs show higher throughput, need less hardware resources and are more energy efficient. In return, fixed-point arithmetic requires to properly deal with overflows and precision losses due to arithmetic round-off errors and less precise value presentation. This loss in precision is manageable for some implementations [249] whereas others require specific handling techniques to ensure the control stability [250]. As an alternative approach, customized floating-point number formats have been explored to fill the large space between the standard double to single-precision formats [84, 235, 251].

The use of alternative number formats also impacts the required memory resources, which is a critical and often limiting factor in MPC implementations on FPGAs. In most related works, the limited memories on FPGAs, i.e. BRAM, limits the realizable system complexity. Techniques like dynamic scheduling have been proposed to deal with limited hardware resources [248]. The authors of [252] e.g. utilized a heterogeneous platform (CPU + FPGA) to involve external memory capabilities. However, only a few papers explicitly have targeted the issue of limited BRAM by proposing specific memory saving techniques [253]. Among those, a promising approach is to exploit the sparsity of the system matrices [254]. Memory savings by a factor of 10 have been reported [255]. Employing compressed matrix formats like *compressed diagonal storage* can lead to substantial memory savings [256], since only matrix diagonals with nonzero entries need to be stored.

[4] Paragraph has already been published in [48]

Depending on the structure and dimension of the matrices memory saving of at least one order of magnitude have been achieved [257]. Utilizing further structural properties in the MPC matrices, reduced the memory requirements by 75% [257].[5]

At the time of this research, most proposals by related work were limited to reference designs or benchmark systems and often only proposed algorithm-specific optimizations. However, only a few works came with a practical relevance to the industrial or environmental infrastructure sector. Therefore, designs that consider specific related properties — e.g., that control systems for environmental infrastructures require long prediction horizons for a stable operation due to their rather slow system dynamic [86] — have been missing in state-of-the-art. In fact, to the best of our knowledge, all proposed works of embedded MPC for FPGAs exhibited a prediction horizon of considerably lower than 100.

[5] Paragraph has already been published in [48]

5.2 New Contributions

This section presents the new scientific contributions of this thesis in the research field of stationary embedded systems, powered by the energy grid. It addresses the design and optimization of embedded MPC controllers on FPGAs in the area of industrial and environmental infrastructures. Such MPC-based controllers are well-known for their computational complexity and memory-intensive algorithms, which leads to the major challenges in the design procedure and implementation of these controllers:

- Section 5.2.1 addresses these design challenges between the high computational requirements of MPC and the restricted computational resources of embedded FPGA-based computing platforms. Specifically, an MPC controller for FPGAs is designed for smart environmental infrastructures, which in this case enables an economic operation of a *water distribution network* (WDN). The controller is optimized to fit into the target device Xilinx Zynq7000, a low-cost FPGA platform that comes as heterogeneous device as it also incorporates a *processing system* (PS) based on an ARM Cortex A8. In order for the controller to fit into the target device, a new optimization method is presented that reduces the memory demand by a factor of ≈ 80 compared to the base design generated by Matlab. This is part of a multidisciplinary project, in which this work covers the optimization and implementation of the controller onto an embedded FPGA-based computing platform. Further contributions are the MPC controller and benchmark design [48, 258], which are documented in the appendix D.2.
- Section 5.2.2 shows how the proposed embedded control system can be physically incorporated into a state-of-the-art pump aggregate.

5.2.1 Structural Data Compression for MPC on FPGAs

The main control objective of environmental systems is to achieve an optimal long-term operation while ensuring a stable service that is robust to external disturbances. Costs like electrical power required for operation are to be minimized while still satisfying the consumer demands. Such systems are usually large scale, subject to constraints and need to be operated both safely and cost-efficiently. In addition, they usually have a slow system dynamic, which requires long prediction horizons for a stable operation [86]. Therefore, design metrics like scalability and cost-efficiency are more important than speed and performance. The reported results are also applicable to other infrastructure systems with these properties, such as irrigation canals, supply chains or traffic control. In this work, the methodology is presented for a system structure derived for a simplified WDN, namely a linearized discrete-time system, subject to box-constraints with time-varying quadratic cost.[6]

The purpose of the control system in this work is to optimize the operation of a WDN system over a long prediction horizon in an economical point of view. A sample rate of 15 minutes was considered, which leads to a prediction horizon N of slightly below 100 in order to cover an entire day.

Basic Controller Design

The mathematical design of the MPC controller setup of this work has been designed in Matlab and is described in the appendix D.2.

There are several approaches to create and implement hardware designs for FPGAs based on existing Matlab reference designs. As discussed in Section 5.1, related work has basically proposed two design flows for MPC that utilize software tools provided by Matlab: First, generating HDL code by using the Matlab *HDL Coder*, and second, generating C-code by using the Matlab *Coder*. The latter approach additional requires an HLS tool to create a hardware out of C-code. Both design flows are shown in Figure 5.3 labeled as `<1>` and **`<2>`**, together with the advanced design flow proposed in this work (**`<3>`**). This advanced design flow is discussed later in this chapter.

[6] Paragraph has already been published in [48]

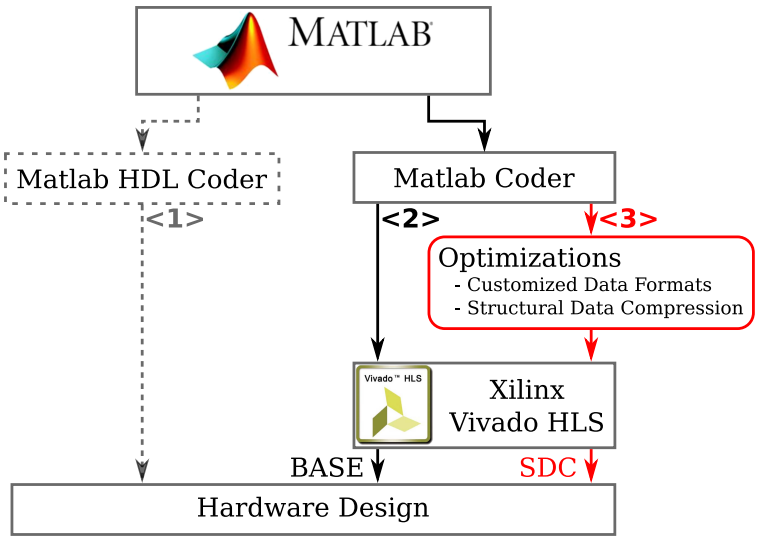

Figure 5.3: Design flows for Matlab-based hardware design.[7]

According to results of related work, in the most cases design flow **<2>** leads to less resource-intensive hardware designs than design flow <1>. Therefore, this approach was used in this work to create a base hardware design out of the existing Matlab reference model. This base design was not synthesizable due to exorbitant memory requirements that exceeded the available memory resources on the Xilinx Zynq7000 target device by a factor of ≈ 30 as shown in Figure 5.4.

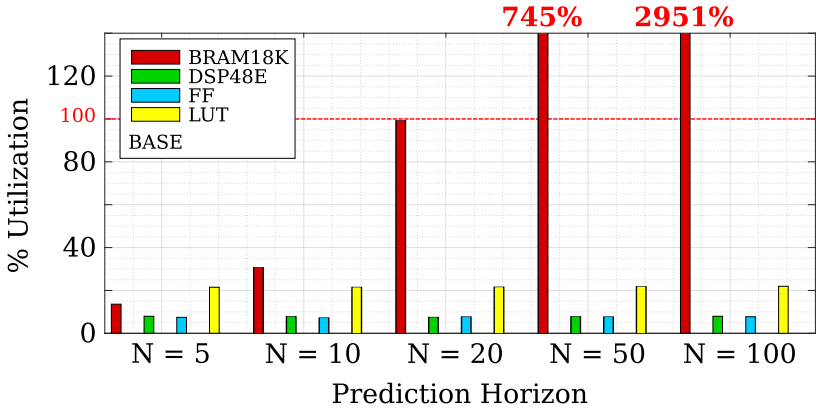

Figure 5.4: Hardware resource utilization of the BASE design.

In order to create a hardware design that fits into the Xilinx Zynq7000 target device, a dedicated optimization strategy was developed. This strategy mainly addressed specific design optimizations to significantly reduce the memory demand. It exploits structural properties in the system matrices, similar to sparse matrix formats. Sparse matrix formats have been utilized by related work to reduce the memory footprint of embedded MPC designs [256, 257]. However, according to the reported outcomes, the required factor of $\approx$

[7] Graphic has already been published in [48]

30 for the system in this work was not achievable with common sparse matrix formats. As a solution, a new strategy was designed, dedicated to the specific system in this work, whose optimization methods lead to a sufficient reduction of the memory footprint. However, these methods incorporated a new memory addressing scheme, which led to an increased demand for logic units that exceeded the resources of the target device. Therefore, a further strategy was developed to optimize this memory addressing scheme. After this, the hardware design was synthesizable for the Xilinx Zynq7000 target device and therewith allowed for further hardware optimizations. Thus, the overall optimization strategy consists of three steps that have been conducted consecutively:

- **Structural Matrix Compression**: Exploiting structural properties of the system matrices such as sparsity and duplications for data compression to reduce the memory demand.
- **Matrix Element Addressing Optimization**: Optimizing the addressing logic for the matrix elements in the compressed matrix memory to reduce the logic cell demand.
- **Hardware Design Optimization**: Optimizing the synthesizable hardware design with respect to performance and energy efficiency by evaluating multiple number formats and incorporating different parallelization techniques.

Structural Matrix Compression

The reason for the large memory requirements for the targeted prediction horizon lies in the relation between the system size and the sizes of the system matrices. Among others, the prediction horizon determines the system size, and, according to Equation D.2.1 and Equation D.2.1, the system matrices $\boldsymbol{H}_t$, $\boldsymbol{A}_{eq}$ and $\boldsymbol{A}_{in}$ as well as the scaling matrices $\boldsymbol{L}_\lambda$ and $\boldsymbol{L}_\mu$ grow quadratically with the system size. Analyzing the structure of these matrices reveals that these matrices are sparse matrices, meaning that the vast majority of entries are zeros. At the algorithmic level, this is done on purpose to simplify the matrix calculations in the optimization algorithm. These calculations follow a straight approach of matrix element addressing and form a clear data flow. Furthermore, for the same reason these matrices contain a lot of duplicates, e.g., the repeated sub matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ inside the $\boldsymbol{A_{eq}}$ matrix as shown in Equation D.2.1.

To reduce the memory footprint, the sparse matrices where partitioned into multiple sub matrices as shown in Figure 5.5. Additionally, masking matrices, i.e., those who only contain identity sub matrices $\boldsymbol{I}$, have been removed and their respective functional aspect has been implemented into the control flow of the algorithm.

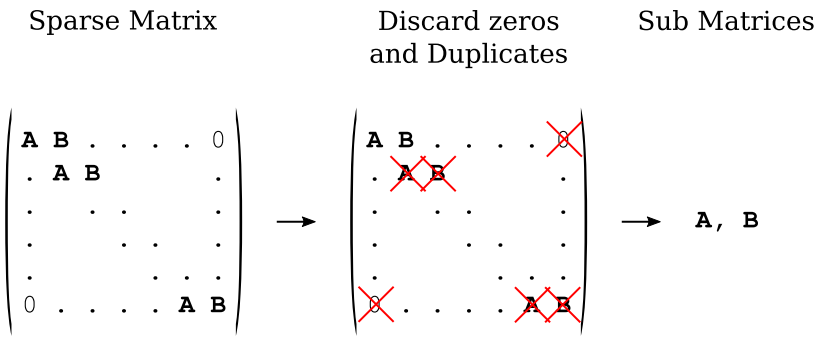

Figure 5.5: Matrix compression by dividing large sparse matrices into multiple sub matrices.

Duplicates in form of identical sub matrices have been removed and the respective access in the algorithm has been changed to point to the respective unique sub matrix. However, this led to issues at the respective locations in the algorithm. Whereas large sparse matrices intentionally contain duplicate elements to simplify the element addressing by straight iteration processes, in the distributed sub matrices without duplicates an additional control flow is required for proper addressing. This additional control flow is required to address the proper sub matrix and to mask specific addresses to point from different addresses to the same element due to the removal of duplicates.

Figure 5.6 shows the synthesis results of the matrix compression for different prediction horizons.

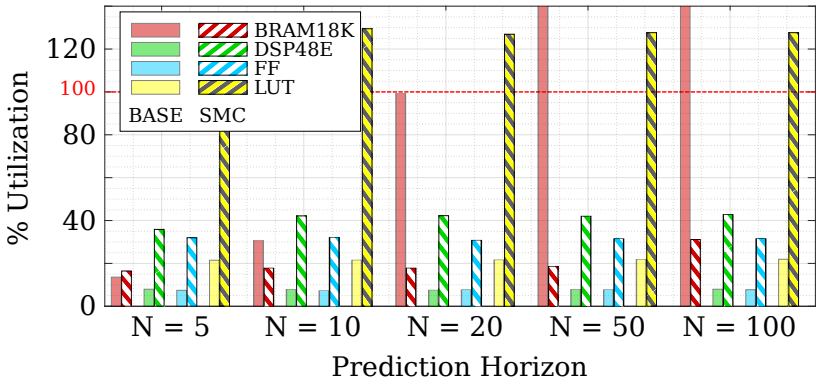

Figure 5.6: Hardware resource utilization after Structural Matrix Compression.

The memory requirements were reduced significantly such that the available memory resources are not a limiting factor any more for the desired prediction horizon of $N = 100$. Particularly, the following system matrices have been replaced in the implementation:[8]

- $\boldsymbol{A_{eq}}$ has been replaced by its sub matrices $\boldsymbol{A}$ and $\boldsymbol{B}$.
- $\boldsymbol{H_t}$ has been replaced by its sub matrices $\boldsymbol{Q_i}$ and $\boldsymbol{R_i}$ for all $i \in [t, ..., t + N - 1]$.
- $\boldsymbol{L_\lambda}$ and $\boldsymbol{L_\mu}$ have been replaced by their diagonal vector.

[8] Bulleted list has already been published in [48]

- Element-masking structures like $\boldsymbol{A_{in}}$ as well as the $\boldsymbol{I}$ components in $\boldsymbol{A_{eq}}$ have been implemented in the control flow of the algorithm.

The significant memory demand reduction with these replacements, however, came at the expense of additional logic resources that is for the additional control flow circuitry. As a result, the LUT utilization increased to a level that exceeded the available resources.

Matrix Element Addressing Optimization

The reason for this increased demand of LUT cells was determined to be the units that are responsible for the matrix element addressing. As described above, these units now incorporate lots of additional control flow due to the matrix compression. Whereas in the base design, the main matrices are constructed to hold copies of their sub matrices at the specific index of the respective iteration and prediction step, new access pattern are necessary to match these different accesses to the same single memory location of the sub matrix. Starting from the different accesses, the element addressing needs to wrap around on the boundaries of the respective sub matrix, i.e., their dimensions. Therefore, the respective control unit for this logic incorporates computational operations like division, rounding and modulo. However, these types of operations are complex to be implemented in hardware, thus requiring a lot of hardware resources. As solution, the dimensions of the sub matrices where extended to match the next consecutive value of $2^n, n \in \mathbb{N}$. This way, the aforementioned computational operations are significantly simplified at the expense of slightly higher memory demands. Figure 5.7 demonstrates this relationship in the synthesis results for different prediction horizons.

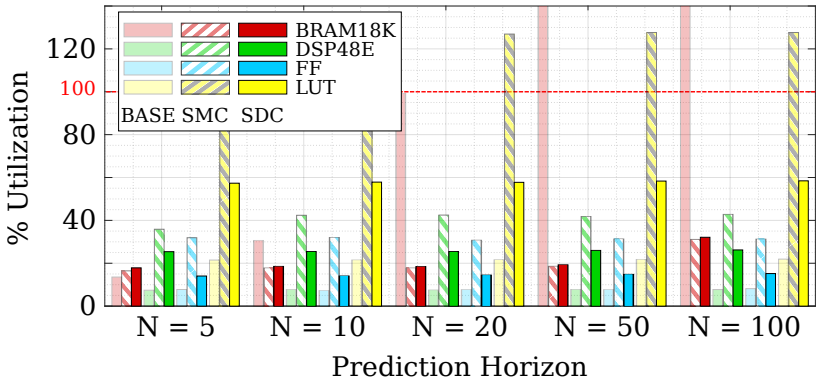

Figure 5.7: Hardware resource utilization after Matrix Element Addressing Optimization.

Hardware Design Optimization and Final Results

The aforementioned design strategies were applied on the algorithmic layer of the MPC design, and led to a synthesizable hardware design that can be implemented on the target

device directly. To that point, however, there is no FPGA-specific optimization in terms of parallelization in the MPC design. Many available hardware resources are not yet utilized by the current design, which makes them available for pipelining strategies. These optimizations promise further achievements towards a faster and more energy-efficient hardware design. Besides parallelization and pipelining, a further advantage of FPGAs is the flexibility of choosing any number format representation.

Number Format

Figure 5.8 shows the hardware resource utilization with respect to different floating and fixed-point number formats.

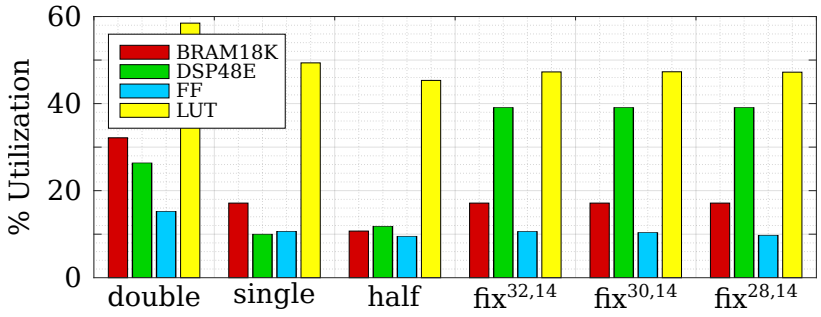

Figure 5.8: Resource utilization of *Structural Data Compression* (SDC) design for different number formats (N = 100).[9]

The choice of the number format is a straight tradeoff between required hardware resources, latency (computation time) and achievable computational precision (also see appendix D.3). The Matlab reference model in general works with double-precision floating-point number format, which leads to the most resource-intensive design in the figure above. However, when changing to a smaller, less precise number format, the computational precision must be considered carefully as it can affect the stability of the control system application. Figure 5.9 shows the latency and the absolute error of the system output with respect to the Matlab reference model for different floating and fixed-point number formats.

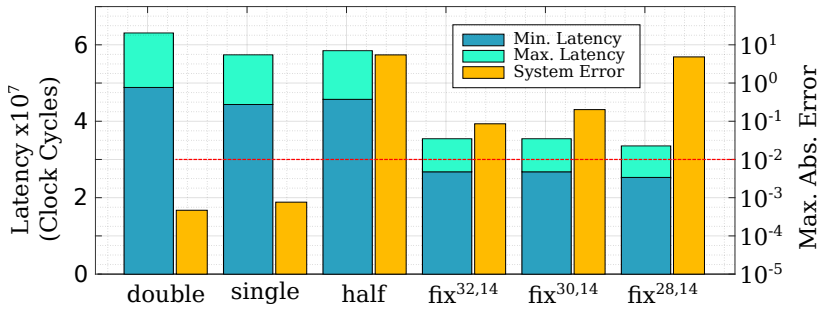

Figure 5.9: Latency and system error of SDC design for different number formats.[9]

[9]Graphics have already been published in [48]

The red-dotted line represents the critical deviation of the system output as it has been determined with the reference model. The stability of the control system cannot be guaranteed if any absolute error of the output exceeds this value.

As shown in Figure 5.9, all employed fixed-point formats as well as the half-precision floating-point format led to a system error that can affect the system stability. Therefore, besides the double only the single-precision floating-point number format was considered in the final hardware design.

Parallelization

Parallel operating design sections and pipelining is realized by inserting directives into the HLS code. As the MPC optimization algorithm in this work follows an iterative approach, the parallelization of specific operations will most likely not lead to significant performance improvements of the system. However, the description of the optimization algorithm, as shown in Algorithm D.1, shows that each iteration step consists of four sub steps, from which step 3 and 4 have no dependencies on each other. The parallelization of the sub steps 3 and 4 only led to a reduction in latency by around 17% as shown in Figure 5.10 and came at no cost of additional hardware resource requirements.

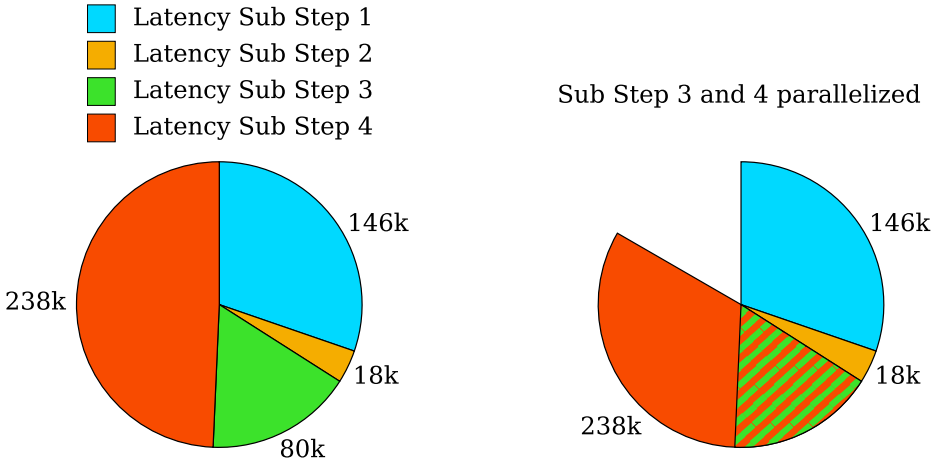

Figure 5.10: Latency of the sub steps of the algorithm and overall latency of a single iteration step for employed parallelization.

In order to improve the performance more significantly, the calculations of the MPC optimization algorithm have been pipelined. As shown in Figure 5.11, this led to a reduction in latency by a factor of ≈ 20 at the cost of additional hardware resources.

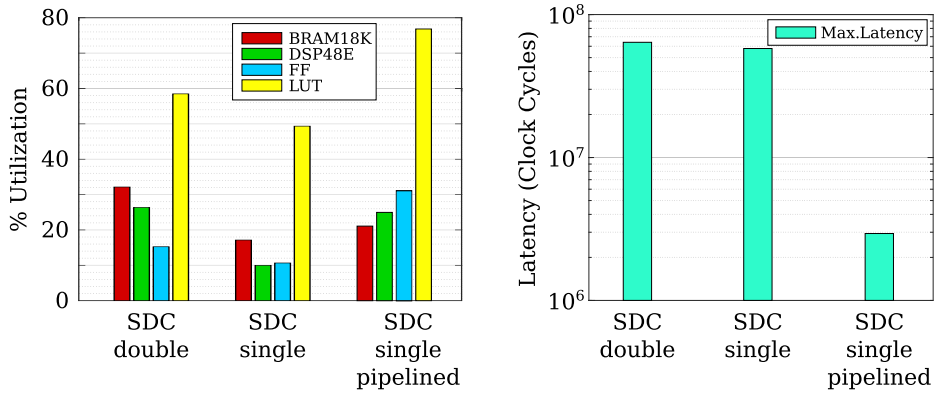

Figure 5.11: Utilization and latency of SDC design ($N = 100$) at a clock frequency of $100MHz$.[10]

Discussion and Conclusion

The proposed SDC techniques lead to a reduced memory demand by factor of ≈ 80 compared to the Matlab-based base design. To the best of our knowledge, there was no optimization strategy reported in related work that yielded a comparably significant reduction in memory demand for embedded MPC at the time of this research.

Furthermore, the proposed work has been the first work that specifically addressed long prediction horizons on embedded FPGA-based MPC. Figure 5.12 shows targeted prediction horizons of related work compared with this work. The presented techniques enable prediction horizons of up to $N = 100$ for double-precision and prediction horizons of up to $N = 500$ for single-precision-based system designs on a low cost FPGA platform.

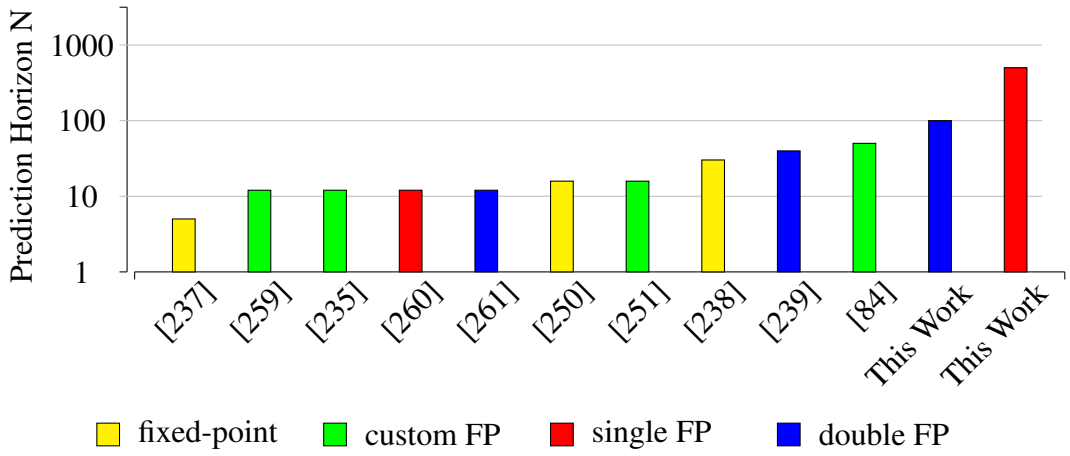

Figure 5.12: Comparison of related MPC designs by implemented prediction horizon and chosen data format like fixed-point or custom, single and double-precision floating point (FP).

[10] Graphic has already been published in [48]

5.2.2 System Emulation

In order to evaluate the proposed control system for real infrastructures, specific methodologies are required to interconnect the control system with physical pump actuators. Therefore, a prototype has been designed that demonstrates the implementation of the controller hardware in an existing pump station sub system. As basis for this prototype, a novel Calio pump aggregate by KSB has been employed. For industrial control purposes, the Calio pump aggregates come with an integrated Modbus [262] interface. In this work, this interface has been used to emulate the integration of the MPC target platform into the pump aggregate as shown in Figure 5.13. This way, no physical changes have to be made on the pump for system emulation.

Figure 5.13: System emulation with Modbus-based communication between Xilinx Zynq computational unit and KSB Calio pump (sources: Xilinx, `www.modbus.org`, `www.ksb.com`).

As evaluation board for proposed Xilinx Zynq *multiprocessor SoC* (MPSoC) hardware platform, a ZedBoard [263] was employed. The Modbus control was implemented on the PS of the Zynq, and its physical layer on the *programmable logic* (PL). An additional interface converter provides a differential signal interface to the Calio pump device. Figure 5.14 shows the system emulation setup in the laboratory environment.

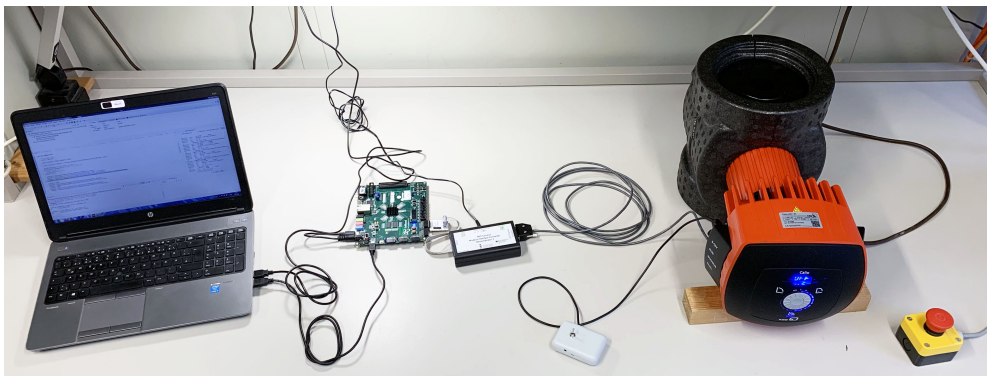

Figure 5.14: System emulation setup (from left to right hand side): Debug laptop, Xilinx Zynq7000 evaluation board, RS-485 interface converter, KSB Calio pump.

This setup exemplarily showed how Calio pump devices can be implemented in experimental control systems setups without changing the hardware of the existing product. This system emulation was provided to the industrial project partners to evaluate the technological possibilities for the proposed concept of embedded high-performance platforms for MPC in decentralized pump system infrastructures.

Chapter 6

Conclusion and Future Work

> ”Es gibt noch viel zu tun, und ich bin sehr beschäftigt.”
>
> Wilhelm Conrad Röntgen

The energy efficiency of microelectronic systems is improving continuously. Since the invention of the first microprocessor, the required energy per bit in compute processing units has continuously been decreasing. However, as the number of deployed microelectronic systems and their computational capabilities is growing faster than ever, the global energy demand for computing is growing exponentially despite the energy efficiency improvements. Additionally, the growing demand for, specifically, mobile electronic devices is causing an enormous and fast growing demand for battery capacity. Therefore, the global production of energy and battery capacity will have to be increased in the near future, which are considerable steps backwards for the efforts towards a more energy-friendly and sustainable future. In order to prevent this increase in energy and battery production, revolutionary changes in microelectronic systems design are inevitable.

This thesis therefore focused on cross-application research topics that help to improve the energy efficiency of stationary and battery-powered microelectronic systems. And, to take yet another step forward, it furthermore faced the latest challenges to replace conventional energy supplies by energy harvesting technologies.

The term *microelectronic system* is more or less a collective term for a vast number of different electronic systems and devices. These systems and devices come in a huge variety of use cases with different requirements, boundary conditions, specific design constrains and design challenges. In order to systematically manage all the different systems and their potential design approaches towards a better energy efficiency, this thesis followed a dedicated research strategy. The first step of this strategy was to divide microelectronic systems into three categories: Energy harvesting-powered autonomous embedded systems, battery-powered mobile embedded systems and energy grid-powered stationary embedded systems. In the second step, yet open scientific challenges and related new design approaches have

been assessed by considering the large design space of different computing architectures, design methodologies and application implementations. The final step was to evaluate these new design approaches on dedicated hardware, i.e., on a real embedded system application.

The tackled scientific challenges and presented scientific contributions in this thesis are summarized as follows for the corresponding categories:

- **Autonomous Embedded Systems** (Energy Harvesting-Powered)

  Critical design challenges are to properly deal with the unforeseeable and sudden power losses, low power output levels of the energy harvesters and the cold-start overheads due to the frequent system restarts. Related work is limited to *reactive* power loss detection principles that do not take harvester-specific characteristics into account. Therefore, they need to react on upcoming power losses immediately, which makes the finalization of atomic operations nearly impossible.

  This work presents new *proactive* power loss detection principles that enable a more harvester-aware transient computing for kinetic and thermoelectric energy harvesters. The presented methods allow for detecting potential power losses up to 10 times earlier than state-of-the-art proposals and thus enable a *proactive* power loss handling like the reliable finalization of atomic operations despite sudden power failures. Another contribution is a new design principle that allows a significant reduction of the system cold-start period to less than 75 milliseconds, which is more than 30 times below the reported values in related work.

  The reported contributions allow for a wider deployment of truly autonomous embedded systems, which only rely on sustainable and renewable energy sources without the need for batteries nor other electrochemical energy buffers.

- **Mobile Embedded Systems** (Battery-Powered)

  Important design challenges for mobile battery-powered devices are maximizing battery lifetimes and improving edge computing capabilities despite a limited energy reservoir. In this work, two application areas with specific design constrains have been considered that comprehensively address these challenges.

  In the field of *wireless sensor networks* (WSNs), the design of wearable sport performance analysis devices has been focused. Here, new principles for low power system design and precise wireless synchronization have been presented that outperform related work. First in terms of usability, as they enable in-field sport performance analysis for athletes in contrast to stationary laboratory measurements. Second in terms of synchronization accuracy, as they show the best synchronization accuracy for low power *Bluetooth Low Energy* (BLE) among state-of-the-art approaches.

  Regarding the emerging application field of *Internet of Things* (IoT), new design principles have been proposed to allow for a low power but high security (*Transport Layer Security* (TLS)-level) wireless communication that tackle the critical aspect of

data security as one of the most crucial challenges in the field of battery-powered IoT edge devices.

- **Stationary Embedded Systems** (Energy Grid-Powered)

  A major design challenge of stationary embedded systems is to realize an energy-efficient implementation of high performance computing tasks. To achieve this, a turning away from conventional sequential computing platforms like *Central Processing Units* (CPUs) towards computing platforms that offer high degrees of parallelism like *Field Programmable Gate Arrays* (FPGAs) is inevitable.

  This work addressed the specific design challenges for a parallelized, low power implementation of *Model Predictive Control* (MPC) on FPGAs. The most critical design constraint of data-intensive computing tasks like MPC is the limited amount of high-bandwidth memory on FPGAs. As a solution, this work proposes *Structural Data Compression* (SDC) as a new method to significantly reduce the memory demand of MPC with long prediction horizons. For the chosen control algorithm, a memory demand reduction by a factor of 80 compared to the reference base design has been achieved, which is higher than the achievements by proposed methods of related work. The underlying application is dedicated to a smart water supply infrastructure system, and thus exemplary shows how the energy efficiency of embedded infrastructure control systems can be improved significantly.

With respect to the second step of the discussed research strategy, Figure 6.1 gives an overview about the new design approaches within the design space in terms of architecture, methodology and application (introduced in Figure 1.6).

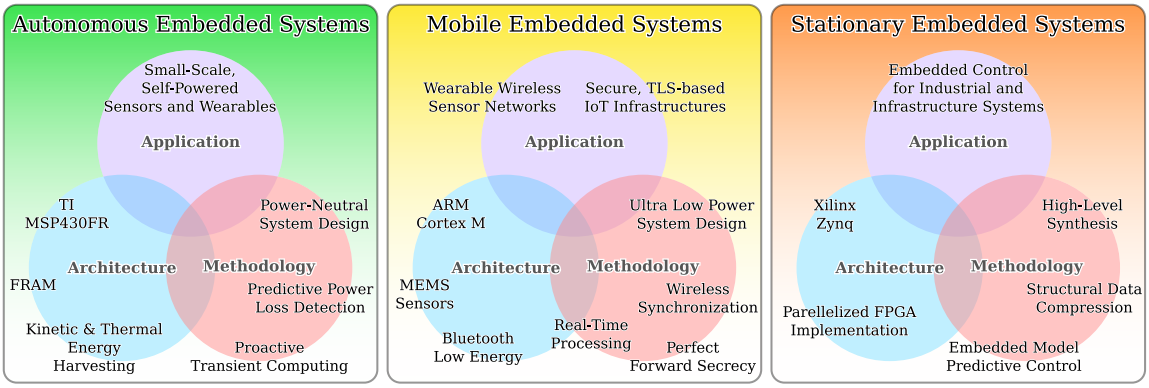

Figure 6.1: Overview about the targeted embedded applications, the applied design methodologies and the employed architectures for the contributions in this work.

Future Work

To once meet the goal of a sustainable global energy utilization, the energy demand of microelectronic systems must be significantly decreased further. Considering that the data-

driven world and its hunger for computing energy is just beginning, it is important to act immediately [3]. In the future, many ongoing developments like smaller technology sizes and non-silicon-based semiconductors will exhibit a better energy performance of computing and thus contribute to less energy-hungry electronics. Emerging quantum computers allow for several computation tasks to be executed significantly faster than on conventional computers and thus conserving energy in the time domain. The increasing number of data-intensive computation tasks, reflected in the growing demand for applications that utilize *machine learning* (ML) and *artificial intelligence* (AI) technologies, leads to the point that the memory-related energy costs become more and more significant for the energy footprint of microelectronic systems. Therefore, emerging memory technologies like *Spin Transfer Torque* (STT)-*Magnetoresistive RAM* (MRAM), that already show low power consumption and high endurance, will be major contributors for a more sustainable energy utilization. Considering that many of these memories are nonvolatile, they hold enormous potential to push the deployment of battery-free energy harvesting-powered devices. As energy harvesting-powered devices are the most promising sustainable alternative to the growing number of battery-powered devices, future work must consider these memory technologies also in this specific context [4].

However, despite all recent improvements in integrated circuit design, conventional chip design is approaching physical limits like never before. The most critical limit is the *Landauer Limit* [264], which states that there is a physical minimum energy level for deleting a bit in digital computing. In other words, the theoretically reachable energy efficiency of computing systems will saturate one day. To illustrate the *Landauer Limit* within the improvements in integrated circuit design during the last decade, Figure 6.2 shows a compute performance forecast of the year 2016 [25]. It indicates that the compute performance that was targeted for the future (blue line) has been reached by 2021 [2].

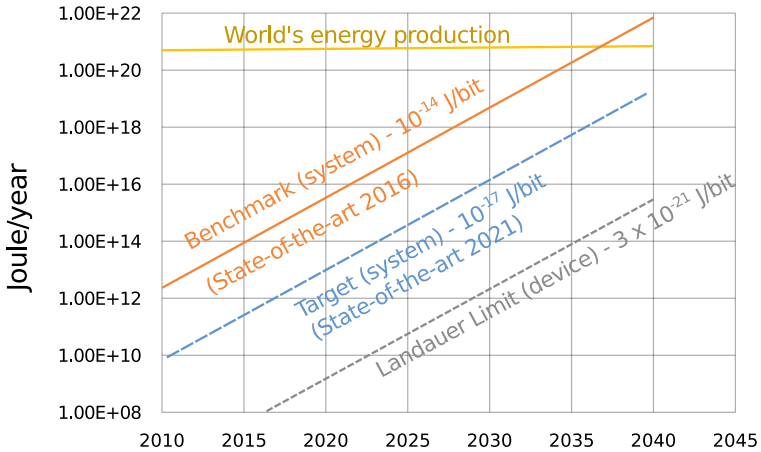

Figure 6.2: Year 2016 forecast about the global energy demand for computing and the energy performance of computing (graphic based on [25]).

But even though it can be said that this goal has more or less been accomplished, it did not contribute sufficiently to suppress the global energy demand for computing [2, 3]. Further-

more, Figure 6.2 emphasizes that even when reaching the *Landauer Limit*, it won't prevent the growing energy demand from approaching the current world's energy production level. Therefore, comprehensive *beyond Landauer* computing paradigms will be inevitable in the future [265] that even might need to replace conventional digital computing [2].

Aside from the microelectronic systems design, a more sustainable energy production holds also much potential to reduce the energy production-related adverse effects on our environment. The same applies to new and pioneering battery technologies that function without critical and seldom materials. As far as state-of-the-art technologies are concerned, fundamentally new and more sustainable manufacturing processes are inevitable. Such as processes that, e.g., consider and foster effective recycling methods especially with respect to battery technologies [30]. Because the current fact is that the growing demand for consumer electronics together with the ever shorter product life cycles caused electronic waste (e-waste) to become one of the fastest-growing waste streams in the world [266]. Thus, apart from any technological aspects, there needs to be a serious rethink on the part of consumers and producers of electronic products. It is indeed remarkable that the current justification for the deployment of new electronic products is their alleged necessity for achieving a more energy-efficient and environmentally friendly world. In reality, however, the related energy demand massively increases despite the promises and, thus, rather contradicts than supports the initial assertion. Moreover, the growing trend of producing and discarding needless products with deliberately short lifetimes and lack of repairability can be seen as the major cause of the dramatically increasing volume of e-waste. Therefore, deploying microelectronic systems truly for the benefit of humanity, and not for purely economic profit or serving temporary fashion trends, will without a doubt be a decisive part of the solution.

Kapitel 7

Zusammenfassung

Die globale Anzahl an aktiv eingesetzten mikroelektronischen Systemen wächst unaufhörlich. Aufgrund kontinuierlicher technischer Weiterentwicklungen werden die Einzelsysteme zunehmend energieeffizienter. Dies reicht jedoch nicht annähernd aus, um den wachsenden globalen Gesamtenergiebedarf aller eingesetzten mikroelektronischen Systeme zu bremsen. Dieser Energiebedarf weist seit Jahren ein exponentielles Wachstum auf, was einer nur linear wachsenden globalen Energieproduktion gegenübersteht. Durch den zunehmend höheren Anteil an mobilen, meist batteriebetriebenen Elektronikgeräten im Konsumerbereich, wächst zudem der Bedarf nach Batteriekapazität exponentiell an. Um einen massiven Ausbau der globalen Energie- und Batterieproduktion mit all ihren fatalen Auswirkungen auf unsere Umwelt und Lebensgrundlagen zu verhindern, sind demnach vielseitige und umfangreiche Fortschritte und Neuentwicklungen notwendig. Im Bereich der Mikroelektronik umfasst dies unter anderem wissenschaftliche Untersuchungen und technische Innovationen im Bereich Hardware- und Systemdesign, mithilfe derer der Energiebedarf stark reduziert oder komplett aus nachhaltiger Energie gedeckt werden kann. Um die rasant steigende Nachfrage nach Batteriekapazität zu bremsen, muss die Verwendung erneuerbarer Energie vor allem bei mobilen Endgeräten im Konsumerbereich vorangetrieben werden.

Behandelte wissenschaftliche Herausforderungen

Die wissenschaftlichen Beiträge dieser Arbeit befassen sich mit verschiedenen Ansätzen, den Energiebedarf mikroelektronischer Systeme einerseits stark reduzieren und andererseits komplett durch Energy Harvesting ersetzen zu können. Die dafür zu bewältigenden wissenschaftlichen Herausforderungen sind eng mit den Eigenschaften der Systeme verknüpft, insbesondere deren Energieversorgung. Deshalb berücksichtigt diese Arbeit drei Kategorien eingebetteter mikroelektronischer Systeme:

- **Autonome eingebettete Systeme** (Energieversorgung über Energy Harvesting)

  Eine Energieversorgung über Energy Harvesting ist nicht beständig und kann jederzeit ausfallen. Solche sogenannten *Transient Computing Systeme* müssen deshalb so gestaltet werden, dass sie unter allen unvorhersehbaren Ausfällen der Energieversorgung zuverlässig funktionieren. Dies birgt viele Herausforderungen beim Systemdesign, die sich grundsätzlich von denen batteriebetriebener Systeme unterscheiden:

  - Sicherstellen eines kontinuierlichen Forward Progresses der Applikation trotz plötzlicher Ausfälle der Energieversorgung sowie das garantierte Abarbeiten atomischer Prozesse.
  - Die typischerweise geringe Ausgangsleistung und kurze Dauer der verfügbaren Energie stellen strikte Randbedingungen an den maximalen Leistungsbedarf und die verfügbare Berechnungszeit dar.
  - Häufige Resets und Neustarts verstärken den negativen Einfluss von Coldstart-Verlusten. Dies kommt insbesondere bei aus Standardkomponenten konstruierten Systemen zum Tragen, da diese Komponenten üblicherweise nicht für diese Betriebsbedingungen konzipiert sind.
  - Die physikalischen Energieumwandlungsprinzipien unterschiedlicher Energy Harvester weisen sehr spezifische Charakteristika hinsichtlich Ausgangsleistung und -dauer auf. Zudem sind die Wechselwirkungen von physikalischer Anregung einerseits und elektrischer Last andererseits stark vom jeweiligen Transduktionsprinzip abhängig. Ein effizientes Systemdesign erfordert deshalb harvesterspezifische Optimierungen und ganzheitliche Systemanalysen.

- **Mobile eingebettete Systeme** (Energieversorgung über Batterien und Akkus)

  Der technologische Fortschritt von Batterietechnologien hinsichtlich Energiedichte kann nicht ansatzweise mit den Fortschritten in der Mikroelektronik hinsichtlich Integrationsdichte und Performanz schritthalten. Dies hat dazu geführt, dass die Batterie zum wichtigsten limitierenden Faktor beim Design tragbarer Gadgets und mobiler *Internet of Things* (IoT) Endgeräte geworden ist. Somit führt die limitiere Batteriekapazität auf der einen Seite und die angestrebte hohe Performanz auf der anderen Seite zur den größten Herausforderungen für die Zukunft des Edge-Computing:

  - Für den jeweils besten Kompromiss zwischen Performanz und Energiebudget sind applikationsspezifische Ultra Low Power Designmethoden erforderlich. Im Bereich tragbarer Sensornetzwerke beispielsweise umfassen diese eine sorgfältige Auswahl der Hardwarekomponenten und ein angepasstes Softwaredesign, um trotz limitiertem Energiebudget ein langzeitliches und präzises Bewegungsmonitoring zu ermöglichen.
  - Im Bereich des IoT wächst die Zahl an batteriebetriebenen Endgeräten unaufhörlich an. Gleichzeitig steigt jedoch auch das Sicherheitsrisiko, da hochentwickelte Sicherheitsmechanismen aufgrund deren berechnungsintensiven kryptografischen Algorithmen und dem damit verbundenen Energiebedarf oft vernachlässigt wurden. So wurde das Thema Security zu einer der größten Herausforderungen im Bereich IoT Edge-Computing.
  - Ein limitiertes Energiebudget limitiert auch die Performanz etwaiger drahtloser Kommunikation. So sind speziell für den Low Power Bereich zugeschnittene Technologien wie *Bluetooth Low Energy* (BLE) gegenüber anderen Verfahren in ihrer Bandbreite und Flexibilität eingeschränkt. Die führt zu großen Herausforderungen bei der Realisierung präziser zeitlicher Synchronisation in tragbaren Sensornetzwerken oder der Implementierung *Transport Layer Security* (TLS)-basierter Sicherheitsmechanismen die zur gegenseitigen Authentifizierung auf die Übertragung großer Zertifikate und Handshake-Nachrichten setzen.

- **Stationäre eingebettete Systeme** (Energieversorgung über das Energienetz)

  Stationäre eingebettete Systeme stellen üblicherweise Subkomponenten anderer Systeme dar. Im Konsumerbereich übernehmen diese bspw. die Steuerung elektrischer Geräte, während sie im industriellen Bereich für komplexe Regelungsaufgaben für industrielle Prozesse eingesetzt werden. Viele industrielle Regelungsaufgaben stellen hohe Anforderungen an Echtzeitfähigkeit und Regelungsstabilität und erfordern zudem die gleichzeitige Berücksichtigung vieler Zustandsgrößen. *Model Predictive Control* (MPC) hat sich hier als Regelungsmethode durchgesetzt, da diese die Anforderungen auch unter strikter Beibehaltung von Randbedingungen erfüllen kann. Diese Methode ist allerdings rechnerisch komplex und für einen hohen Energieaufwand bekannt, da sie auf einer kontinuierlichen mathematischen Optimierung der Regelgröße innerhalb eines bestimmten Prädiktionshorizontes basiert. Deshalb sind große Herausforderungen für eine eingebettete Implementierung auf energiebeschränkten Systemen zu meistern:

  - Anspruchsvolle Berechnungsverfahren wie MPC können von Hardwareplattformen wie *Field Programmable Gate Arrays* (FPGAs), die einen hohen Parallelisierungsgrad ermöglichen, profitieren. Verglichen mit sequentiellen Berechnungen auf *Central Processing Units* (CPUs), ermöglichen dedizierte Hardware-

strukturen auf FPGAs vergleichbare Performanzen bei weitaus geringeren Taktfrequenzen und Energiebedarfen. Allerdings birgt das Systemdesign aufgrund FPGA-spezifischer Charakteristiken große Herausforderungen, wie bspw. die limitierte Verfügbarkeit von Speicher- und Hardwareressourcen.

- Dies führt zu Design Tradeoffs zwischen Speicherbandbreite, Parallelisierungsgrad, erforderlichen Hardwareressourcen und Echtzeitfähigkeit.
- Hardwareimplementierungen erfordern einen anspruchsvolleren und signifikant höheren Implementierungsaufwand als Softwareimplementierungen. Die Energieeffizienz steigt zudem mit zunehmender Spezifizierung auf die entsprechende Anwendung, sodass energieeffiziente Systeme immer mit einem Verlust an Flexibilität einhergehen.

Wissenschaftliche Beiträge

Die meisten in dieser Arbeit vorgestellten Beiträge wurden in Konferenz- und Journalpublikationen veröffentlicht [31–59]. Viele dieser Beiträge wurden in Kooperation mit Industriepartnern und anderen Wissenschaftlern erarbeitet. Diese Personen und Partner sind zusammen mit den wissenschaftlichen Beiträgen folgend aufgeführt:

- **Kapitel 2:** ***Proaktives Transient Computing für eingebettete Energy Harvesting Systeme*** [47, 51, 58]

  Der state-of-the-art Ansatz zur Sicherstellung eines kontinuierlichen Forward Progresses der Applikation trotz plötzlicher Ausfälle der Energieversorgung ist unter dem Begriff Checkpointing bekannt [60–62]. Hierbei wird der flüchtige Systemkontext des Prozessors in nichtflüchtigen Speichern gesichert, und kann so einen Ausfall der Energieversorgung überdauern. Jedoch sind bisher publizierte Ansätze rein reaktiver Natur und können Ausfälle der Energieversorgung nur kurzfristig detektieren. Dies bringt vielseitige Probleme mit sich, in dem es bspw. schwierig bis unmöglich ist, die Abarbeitung atomischer Prozesse sicherstellen zu können.

  In diesem Kapitel wird ein neues proaktives und harvesterorientiertes Verfahren präsentiert, das den internen Zustand von kinetischen Harvestern berücksichtigt um potentielle Stromausfälle voraussagen zu können, was wiederum eine zuverlässige Ausführung atomischer Operationen ermöglicht. Weiterhin wird eine neue Methode für prädiktives Checkpointing vorgestellt, bei welcher der einen thermoelektrischen Harvester anregende Temperaturgradient gemessen wird, um den zukünftigen Verlauf des verfüg-baren Energiebudgets bestimmen zu können. Damit wird ein proaktives Handling von Stromausfällen auf Transient Computing Systemen ermöglicht, da potentielle Stromausfälle bis zu einer Größenordnung früher als in Related Work erkannt werden können. Zudem werden spezielle Optimierungen vorgestellt, mit denen die Kaltstartperioden von Transient Computing Systemen verkürzt werden. Dies ist

essentiell für powerneutrale Systeme deren Harvester nur kurze Ausgangsleistungssignale erzeugen kann, da in diesen Fällen die Länge der Kaltstartperiode darüber entscheidet, ob das System überhaupt funktionieren kann. Während Related Work die Kaltstartperioden in ihren Arbeiten entweder ignoriert oder mit Werten im Sekundenbereich beziffert hat [63], zeigen die Ansätze in dieser Arbeit wie Kaltstartperioden auf bis zu 75 Millisekunden minimiert werden können.

Teile der Arbeiten wurden zusammen mit M.Sc. Vitor Kieling erarbeitet und sind in seiner Masterthesis dokumentiert [87].

- **Kapitel 3:** ***Tragbare drahtlose Sensor Netzwerke für Leistungsanalysen im Profisport*** [32–36, 39–41, 46]

  Hohe Batterielaufzeiten mobiler mikroelektronischer Geräte können nur durch eine umfassende Kombination von eingebetteter Hardware und Software erreicht werden (Hardware–Software Codesign). Dazu basiert das in dieser Arbeit präsentierte Systemdesign auf einem anwendungsspezifischen Cross-Layer Ansatz von Software und Hardware. Um möglichst viele der entsprechenden Herausforderungen in Rahmen eines Anwendungsbereichs behandeln zu können, ist dieses Kapitel dem Entwurf mobiler und tragbarer Sensornetzwerke für den Profisportbereich gewidmet, in welchem elektronische Sensoren für die sportliche Leistungsbewertung im Feld zunehmend an Bedeutung gewinnen [64, 65]. Es wurde gezeigt, dass sich kleinformatige Sensoren wie *micro electromechanical systems* (MEMSs) und *Inertial Measurement Units* (IMUs) für das Erfassen und Analysieren der Bewegungen des Sportlers besonders eignen [66–68]. Das elektronische Design solcher tragbaren Sensoren birgt jedoch große Herausforderungen, aufgrund des limitierten Energiebudgets und der eingeschränkten Rechenkapazität eingebetteter Geräte [69].

  In diesem Kapitel werden daher neue Konzepte und Methodiken für drahtlose und tragbare Sensoren präsentiert, deren Performanz und Messgenauigkeit sich für Anwendungsfälle der Leistungsanalyse im Profisport mit kommerziellen Laborlösungen vergleichen lässt. Zudem wird ein neues Verfahren zur drahtlosen Zeitsynchronisation über den Funkstandard BLE vorgestellt, das ein Echtzeitmonitoring trotz Ultra Low Power Anforderungen auf mehrere Sensoren verteilt ermöglicht. Dieses Verfahren nutzt das charakteristische Stromprofil des BLE Chips während der Funkaktivität aus und kommt auf einen statistischen Synchronisationsfehler von weniger als 10 Mikrosekunden, was eine höhere Genauigkeit als sämtliche bisherigen Verfahren von Related Work [70–73] darstellt.

  Die Arbeiten wurden in Kooperation mit dem Institut für Sport und Sportwissenschaft der Technischen Universität Dortmund und dem Lehrstuhl von Prof. Dr. Thomas Jaitner durchgeführt. Entsprechende Forschungsergebnisse sind in der Dissertation von Dr. Marcus Schmidt [88] und werden in der zurzeit in Arbeit befindlichen Dissertation von Dipl.-Ing. Sebastian Wille behandelt.

- **Kapitel 4:** ***Low Power Highend Kommunikationssicherheit für das Internet der Dinge*** [52–55, 59]

  Ausgereifte Protokolle zur Kommunikationssicherheit wie TLS basieren auf komplexen kryptografischen Algorithmen, die berechnungsintensiv und bekannt sind für ihren damit verbundenen Energiebedarf [74–77]. Aus diesem Grund wurden im Bereich der Related Work viele leichtgewichtige Protokolle als energiesparsame Alternative zu TLS für energie- und bandbreitenlimitierte IoT Endgeräte veröffentlicht [78–81]. Allerdings kommen diese Alternativen grundsätzlich mit Einbußen im Securitylevel daher, sodass viele state-of-the-art Ansätze im Endeffekt ein erhöhtes Sicherheitsrisiko für Edge-Computing darstellen.

  In diesem Kapitel werden neue Ansätze vorgestellt, mittels derer eine Realisierung von rechenleistungsintensiven Protokollen wie TLS auch auf Endgeräten mit stark limitiertem Energiebudget möglich wird und gleichzeitig eine Batterielaufzeit von mehreren Jahren garantiert werden kann.

  Die Untersuchungen wurden in Kooperation mit der KSB AG in Frankenthal und anderen Wissenschaftlern des Lehrstuhls Entwurf Mikroelektronischer Systeme der Technischen Universität Kaiserslautern durchgeführt. Zugehörige Forschungsarbeit wird in den zurzeit in Arbeit befindlichen Dissertationen von M.Sc. Frederik Lauer und M.Sc. Maximilian Schöffel thematisiert werden.

- **Kapitel 5:** ***Eingebettete modellprädiktive Regelung für technische Infrastruktureinrichtungen*** [48]

  Für große Systeme stellt MPC die adäquate Methode innerhalb der Regelungstechnik dar, da gegenseitige Abhängigkeiten mittels mathematischer Vorhersage und Berücksichtigung von möglichen zukünftigen Zuständen in die Berechnung der optimalen Regelungsaufgabe einfließen (Prädiktionshorizont). Der Fortschritt bei neuartigen Berechnungsplattformen mit hohem Parallelisierungsgrad wie bspw. FPGAs, ermöglichte im vergangenen Jahrzehnt auch eingebettete Implementierungen von MPC [82–85]. Allerdings waren die Möglichkeiten hinsichtlich Systemkomplexität und Länge des Prädiktionshorizonts aufgrund der limitierten Hardwareressourcen - wie bspw. Speichereinheiten - zum Zeitpunkt dieser Forschungstätigkeit stark eingeschränkt. Aus diesem Grund waren FPGA basierte MPC Systeme für smarte Infrastrukturlösungen, bei denen ein langer Prädiktionshorizont zur Stabilität des Regelungssystems erforderlich ist, mehr oder weniger nicht einsetzbar [86].

  Im Rahmen dieses Kapitels wird eine neue Strategie zur Datenkomprimierung für eingebettetes MPC auf FPGAs vorgestellt, mit der ein signifikant längerer Prädiktionshorizont als in Related Work möglich wird. Die Regelungsaufgabe ist hierbei auf die Regelung eines intelligenten Wasserversorgungssystems ausgerichtet, deren Low Power Implementierung in die Elektronik einer modernen Wasserpumpe eingebettet werden kann.

Die Untersuchungen wurden vom Bundesministerium für Wirtschaft und Energie (BMWi) im Rahmen des Projektes NET-Control (Förderungsnummer 03ET1278C) gefördert. Dieses Projekt wurde in Kooperation mit dem Lehrstuhl für Regelungssysteme von Prof. Dr.-Ing. Steven Liu an der Technischen Universität Kaiserslautern bearbeitet. Die zugehörigen wissenschaftlichen Untersuchungen im Bereich der Regelungstechnik werden in der Dissertation von Dr.-Ing. Felix Berkel [89] behandelt.

Neben den vorgestellten wissenschaftlichen Beiträgen wurden weitere Arbeiten publiziert, die nicht explizit in dieser Arbeit behandelt wurden:

- Analyse des Retentionverhaltens von *Double Data Rate* (DDR)3 und DDR4 *Dynamic Random Access Memory* (DRAM) *Dual Inline Memory Modules* (DIMMs) [38, 43–45].
- Kapazitives Monitoring menschlicher Nahrungsaufnahme [31].
- Neue Methoden zur Wissensvermittlung im Physikunterricht mit Hilfe von Augmented Reality [50, 56].
- Neue Implementierungsverfahren zu low power Machine Learning Methoden auf eingebetteten *Graphics Processing Units* (GPUs) [49].
- Neue Systemkonzepte für präzise *radio frequency* (RF) basierte Indoor Lokalisierungssysteme [37, 42].
- Neue Methoden zur Integration miniaturisierter Sensorik in Maschinenelemente zur Überwachung von Betriebszustand und Schmierstoffqualität [57].

Appendix

Appendix A

Developments in Integrated Circuits Design

A.1 Developments in Mixed-Signal Semiconductors

Besides digital semiconductors, also mixed-signal *integrated circuits* (ICs) have shown a huge success. Prominent representatives are *radio frequency* (RF) radio chips and modules, whose emergence significantly simplified the challenging RF transceiver design and therewith paved the road for massive deployment of wireless connectivity and the rise of *wireless sensor networks* (WSNs) and the *Internet of Things* (IoT) [267].

In 1999, Silicon Labs released one of the first RF ICs, squeezing most components required for radio communication onto one single chip [268]. This development can be seen as the enabler for massive deployment of wireless connectivity in, e.g., modern cell phones. The manufacturing emerged tremendously since then, such that nowadays, many *Microcontroller Unit* (MCU) manufacturers incorporate such RF circuits on their chips and provide RF-*System on Chip* (SoC).

But the relatively high energy demand of wireless communication was a challenge for battery-powered devices for a long time. In 2010, it was faced by the release of the new *Personal Area Network* (PAN) technology *Bluetooth Low Energy* (BLE) [269], which put the focus on minimizing the energy demand instead of increasing the communication bandwidth like other wireless technologies such as *wireless fidelity (WLAN)* (WiFi) [270]. Half a year later, Nordic Semiconductor released one of the first BLE chips [271, 272]. For the first time, it was possible for battery-powered devices to provide a more or less continuous wireless connectivity and maintaining years of operation. Thus, mixed-signal RF-ICs and BLE can be seen as two of the enabling key technologies for battery-powered IoT edge devices.

A.2 Developments in MEMS

Another technology that has evolved in the wake of the semiconductor progress are *micro electromechanical systems* (MEMSs). MEMS contain movable micro-scale structures like springs, cantilevers and moving plates (Figure A.1). These structures allow to form capacitive or ohmic sensing elements for environmental forces like acceleration, pressure and sonic waves or even electrostatically controlled actuators for optical beam forming.

In the 1980s, silicon emerged as the base material for MEMS, enabling to manufacture the sensing elements and signal-processing circuitry on the same substrate [273]. Early commercially available MEMS sensors targeted automotive purposes [274]. In 1993, the first massively deployed MEMS accelerometer by Analog Devices boosted the deployment of airbags in cars, as it was smaller, cheaper, more energy-friendly and more reliable than conventional technologies at that time [275].

Meanwhile, the MEMS market has almost doubled within the last decade and made up almost 3% of the semiconductor market in 2020 [276]. Today, MEMS has evolved as the state-of-the-art technology for inertial and environmental sensors, audio microphones and transducers, optical and photonic found in all smartphones and IoT sensors and they incorporate more and more dedicated signal processing features up to embedded *machine learning* (ML) and *artificial intelligence* (AI) [276]. In fact, only the progress MEMS enabled the realization of low-energy, cheap and small sensing, which many modern technologies and products rely on. Similar to semiconductor progress, the structures become smaller and smaller due improving production processes, leading towards *Nanoelectromechanical system* (NEMS) [273].

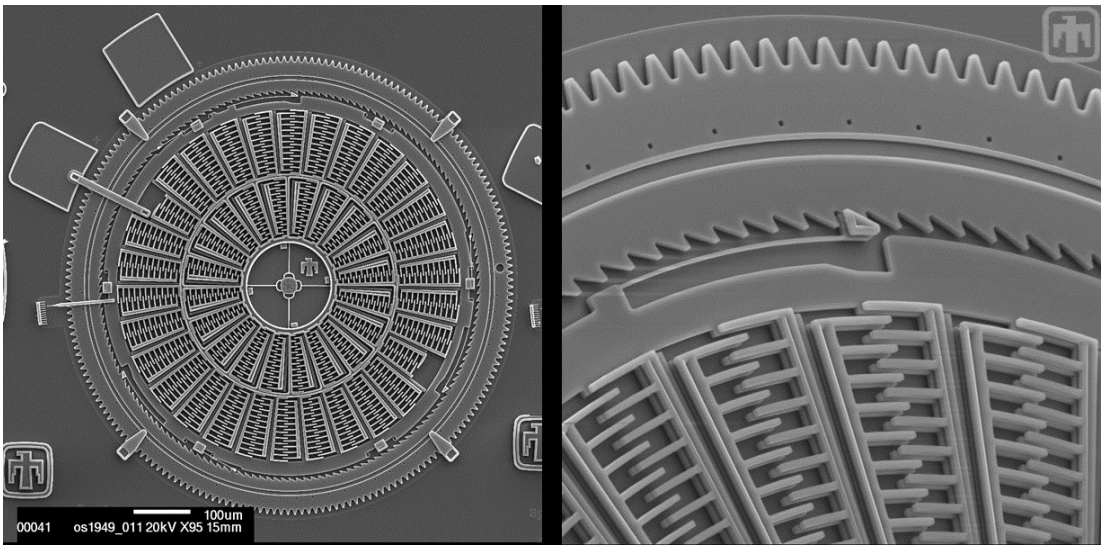

Figure A.1: Example of a MEMS layout: A torsional ratcheting actuator [277].

Appendix B

Embedded Transient Computing Systems with Energy Harvesting

B.1 The *EternAlyzer* Evaluation Platform

The *EternAlyzer* is a transient computing hardware device that has been designed as the evaluation platform for the investigations of transient computing systems in this thesis, specifically those discussed in Section 2.2.1 and Section 2.2.3. It is designed as a wearable gadget that is powered by a micro kinetic energy harvester, the MGS26.4 by Kinetron [164]. This harvester, as discussed in Section B.2.1, is specifically designed for wrist-worn devices to scavenge energy from human motion. Figure B.1 gives an overview of electronic components of the *EternAlyzer*, such as a voltage-doubler rectifier, the buck/boost DC/DC voltage converter LTC3129, the MSP430FR5869 *Microcontroller Unit* (MCU), a memory-*liquid crystal display* (LCD) and the ADXL362 digital *micro electromechanical system* (MEMS) accelerometer. The circuitry from the harvester to the voltage converter is denoted as supply voltage conditioning unit, the MCU and its periphery as load system.

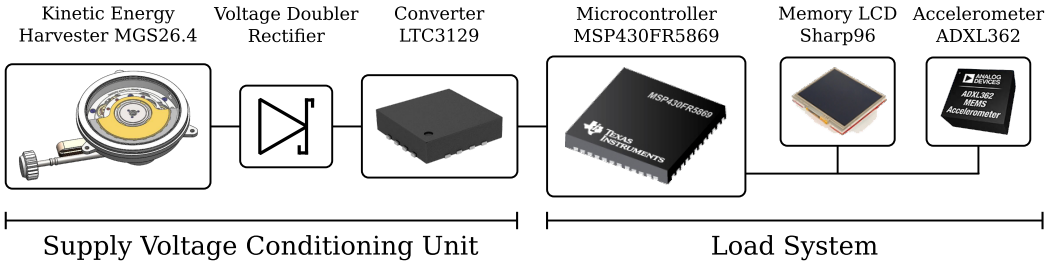

Figure B.1: Hardware components of the *EternAlyzer*.

Its circuitry follows a *Power-Neutral*, battery-less system design that is solely and directly powered by the energy harvester. The *Power-Neutral* principle is mostly reflected in the supply voltage conditioning unit, which represents the major difference to conventional,

battery-powered system design. As discussed in Section B.2.1, human motion-excited kinetic harvesters typically emit the harvested energy in short duration output pulses. Therefore, the *EternAlyzer*, and in particular its supply voltage conditioning unit, has been specifically designed to exhibit low cold-start times. Figure B.2 shows the block schematic of this unit.

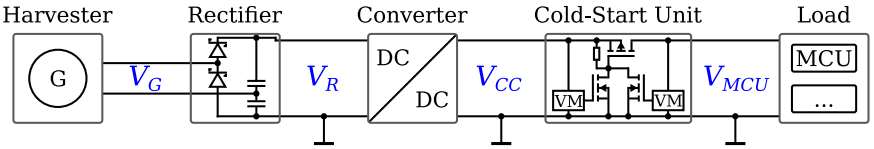

Figure B.2: Supply voltage conditioning unit of the *EternAlyzer*.

V_G is the raw *alternating current* (AC) output voltage of the harvester G that is rectified (V_R) and converted by the DC/DC converter to a stabilized supply voltage V_{CC}. The *Cold-Start Unit* (CSU) — proposed by Balsamo et al. [94] — ensures that the supply voltage to the MCU (V_{MCU}) is not enabled before the minimum operating voltage is reached to allow for a reliable system start-up [94]. Therefore, the CSU incorporates two voltage monitors (VM) with different output voltages to provide a switching hysteresis. For the cold-start investigations discussed in Section 2.2.3, this CSU has been bypassed. Figure B.3 shows the assembled *EternAlyzer*.

Figure B.3: Bottom view on the harvester and top view on the display of the assembled *EternAlyzer*.

This *EternAlyzer* evaluation platform has been chosen as its system design represents the scientific challenges for *Power-Neutral* transient computing systems addressed in this work: The micro kinetic harvester only offers a low power output, which moreover comes at short output durations. This results in many sudden power losses together with many system restarts, which places high challenges on the power loss and cold-start handling. Further challenges are involved by the design of a harvester-specific and comprehensive voltage conditioning, which is inevitable for achieving a high overall system energy efficiency. Moreover, as the *EternAlyzer* is designed as a battery-less and purely energy harvesting-powered wearable device, it illustrates how an environmentally friendly alternative can be realized for the popular sector of wearable electronic gadgets. To the best of our knowledge, at the time of its development in the year 2017, the *EternAlyzer* has been the first *Power-Neutral* system design that comes as a wearable device.

B.2 Small-Scale Energy Harvesting Technologies

This section gives an overview about common energy harvesting principles. The focus is on small-scale harvesters that are — in terms of dimension — suitable to power microelectronic devices, *Internet of Things* (IoT) sensors and wearable gadgets.

B.2.1 Kinetic Energy Harvesting

Kinetic energy harvesters convert mechanical kinetic energy into electrical energy. The most common transduction principle is electromagnetic induction, where, e.g., a specific component that holds a permanent magnet, is set in motion by the kinetic energy and induces a change of a magnetic flux Φ in a coil. This changing magnetic flux induces an electrical voltage in the coil. Small-scale kinetic harvesters that are driven by human motion have been popular in the watch industry for many decades [278]. Nowadays, they are valued as one of the most promising energy harvesting technologies promising for wearable devices [99]. Figure B.4 show the MGS26.4 by Kinetron [164], which has been used in this work.

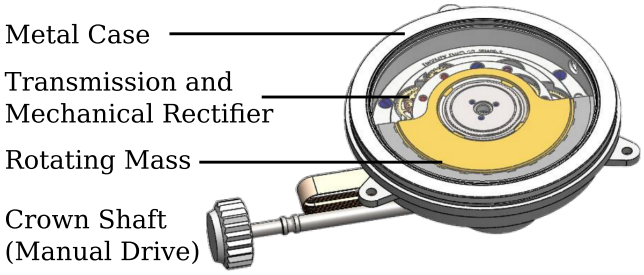

Figure B.4: Micro Kinetic Harvester MGS26.4 [164].

The electrical output voltage of small-scale human motion-driven electromagnetic energy harvester is AC with a couple volts. It requires a rectifier and a fast voltage conditioning circuit, as the output duration is usually only a couple hundred milliseconds per excitation.

B.2.2 Thermal Energy Harvesting

Typical thermal energy harvesters are *thermoelectric generators* (TEGs) or Peltier elements. Their transduction principle is the thermoelectric effect, also known as the Seebeck effect. This effect describes that a temperature gradient, which is applied to two different interconnected metal conductors, leads to a voltage between the two ends of the conductors. State-of-the-art TEGs have two plain surfaces, between which semiconductor material is sandwiched and electrically connected in series. The output voltage V_{TEG} depends on

the applied temperature gradient ΔT, the Seebeck constants of the incorporated materials, and the wiring of the internal structure.[1]

Figure B.5 show a *Prometheus* module by Matrix Industries [165], which has been used in this work. This module consists of a TEG and a voltage boost converter.

Figure B.5: Small-scale thermoelectric harvester module *Prometheus* [165].

The electrical output voltage of small-scale TEGs is *direct current* (DC) with a couple hundred millivolts. For an efficient operation a voltage boost converter with *maximum power point tracking* (MPPT) is required.

B.2.3 Solar Energy Harvesting

Solar energy harvesting is a very popular energy harvesting method due to its broad appliance from tiny solar-powered calculators over autonomous power supply for satellites to photovoltaic power plants. The most common harvesting devices are solar cells that are made out of semiconductor material. Solar cells exploit the photoelectric effect, in which photons cause the movement of load carriers in a semiconductor. The photoelectric effect in general describes the physical interaction of photons in matter and therefore covers various physical effects. Common to all effects is the dissolving of electrons out of their bonding due to being hit by photons, where the photons loose part of their energy (Compton effect).

Commercial solar cells are mostly manufactured out of silicon, but there are also organic solar cells, cheaper, but up to now lower efficiency and lower lifetime. Depending on the thickness of the cell, they are classified as thick film or thin film. Another classification is based on the atomic crystal structure of the silicon, such as mono crystalline and poly crystalline. Amorphous cells, in contrast, do not have a crystal structure.

[1] This paragraph has already partially been published in [58]

The electrical output voltage of small-scale solar cells is DC with a couple volts. For an efficient operation an MPPT is required as their *maximum power point* (MPP) varies with the intensity of solar irradiation.

B.2.4 Radio Frequency Energy Harvesting

Typical harvester devices to transform electromagnetic fields to electrical energy are antennas and coils. The energy output is usually very low, since the field strength and thus the power density of *radio frequency* (RF) fields decay quadratically in their intensity with their distance to the RF emitter. Thus, RF harvesting is mostly utilized in near field to the transmitter, such as for *radio frequency identification* (RFID) transponders or *Near Field Communication* (NFC) systems, or require strong transmitters as in wireless charging applications.

B.3 Nonvolatile Memory Technologies for Transient Computing Systems

In contrast to conventional embedded systems, transient computing systems in general require *nonvolatile memory* (NVM) that can be written to at run-time in order to employ checkpointing. The most common NVM technology is Flash memory, which is also often employed as program memory of MCUs. Flash is one of the best NVM technologies in terms of density and scalability [279, 280]. However, when considered for transient computing systems Flash comes with serious drawbacks. The limited endurance can limit the lifetime of the entire system when exhibiting frequent power losses. Furthermore, Flash exhibits a relatively high write energy demand.

Alternative technologies have been applied for transient computing systems and shown to outperform Flash in case of endurance and write energy demand. Table B.1 gives an overview of these technologies and shows a comparison based on the key characteristics that are relevant for transient computing systems.[2] Further technologies are subject of current research, but in this section only those are discussed that are commercially available.

Technology	Program time per cell	Program energy per cell	Endurance	Available capacity per device
Flash (NAND)	$1...10\mu s$ [281] + $1ms$ for erase	$1nJ$ [281] + $1nJ$ for erase	10^5 [156]	$8TBit$ [282]
FRAM	$50ns$ [283]		$10^{13}...10^{15}$ [156, 283, 284]	$8MBit$ [284–286]
RRAM	$10ns$ [156]		10^5 [287]	$512kBit$ [287]
CBRAM	$10...100ns$ [281]	$0.1...10pJ$ [281]	10^4 [288]	$512kBit$ [288]
MRAM	$12ns$ [156]	$1nJ$ [156]	∞	$32MBit$ [289, 290]
STT-MRAM	$2...20ns$ [291]		∞	$1GBit$ [292]

Table B.1: Comparison of different nonvolatile memory technologies.

B.3.1 Ferroelectric RAM (FRAM)

A *ferro-electric RAM* (FRAM) cell basically consists of a transistor and a ferroelectric capacitor, where the capacitor is the storage element. Its dielectric contains a ferroelectric material that has a crystal structure, in which an atom can be moved between two

[2] Note: This table does not fulfill the requirements of a scientific memory comparison, it rather intends to give an overview about out-of-the-box technologies and devices that are commercially available (by 2020).

stable positions in the lattice by applying an electric field. The position of the atom determines the polarization of the dielectric and therefore has a major effect on the capacity of the capacitor. The most common ferroelectric material employed in FRAMs is Lead Zirconate Titanate (PZT). FRAM chips are commercially available from various manufacturers, who also provide further information about this technology: Texas Instruments [283], Cypress [293], Fujitsu [286] and Lapis Tech [285]. With the MSP430FR series, Texas Instruments offers a 16Bit MCU with embedded NVM as unified memory.

The advantages of FRAM compared to other NVM technologies are the nearly unlimited endurance and its high energy efficiency. Disadvantages are the bad scalability and its limited operation speed.

B.3.2 Resistive RAM (RRAM)

Resistive RAM (RRAM) cells basically consist of a transistor and a resistor, whereas the resistor is the storage element. Logical zeros and ones are stored as specific levels of conduction of the resistor that are established by electrochemical processes. A commercially available RRAM variant that comes with enhanced ultra-low power features is Conductive Bridging RAM (CBRAM) [281]. CBRAM has shown to withstand harsh environmental conditions and keeps cell retention at high temperatures [294, 295]. With the MN101L series, Panasonic offers an 8Bit MCU with embedded RRAM [287]. More information about this technology is provided by one further manufacturer Adesto & Dialog Semiconductor [288].

The advantages of RRAM compared to other NVM technologies are the fast operation speed, its high energy efficiency and its cheap manufacturing process [291]. The major disadvantage is the limited endurance, which makes it unattractive for transient computing systems.

B.3.3 Magnetoresistive RAM (MRAM)

Magnetoresistive RAM (MRAM) cells basically consist of a transistor and a magnetic tunnel junction (MTJ) element, whereas the MTJ is the storing element. A MTJ consists of two magnetic layers and a dielectric tunnel barrier in between. One of the magnetic layers has a fixed magnetic polarization whereas the polarization of the other layer can be set by electric currents. The resistance of the MTJ is determined by the polarization of the free layer, since the magnetic fields qualify how electrons can pass through the barrier by tunneling. More information about this technology is provided by the manufacturers Renesas [289] and Everspin [290, 292]. An advanced MRAM technology is Spin-transfer Torque MRAM (STT-MRAM). STT-MRAM uses spin-polarized currents and thus allow significantly lower energy demands and higher scalability [291].

The advantages of MRAM compared to other NVM technologies are the basically unlimited endurance and its high operation speed. Disadvantages are the worse scalability and its limited energy efficiency due to high write currents. The advanced STT-MRAM in contrast come with a good scalability and high energy efficiency.

Appendix C

Low Power Wireless Communication

C.1 Basic Concepts of Bluetooth Low Energy

Bluetooth Low Energy (BLE), also known as *Bluetooth Smart*, is a wireless communication technology, optimized for low power appliances. In battery-powered devices it facilitates lifetimes up to many years. With common smart phones supporting this technology, its application area is very large. BLE is based on the client-server architecture. Depending on the functional point of view, the client device is denoted as central or master, the server as peripheral or slave. To save energy, BLE is founded on a noncontinuous connection principle. Unconnected slaves frequently announce themselves by broadcasting packets (called advertising). Every master in range that is listening to slaves (called observing), receives these packets. In order to establish a connection, the master sends a connection request to the slave. The connection request packet contains several connection parameters. One of these parameters is called *connection interval* (CI). The CI is a time period mostly in the tens of milliseconds. It determines how frequently from now on master and slave wirelessly communicate. Each CI period is started by a *connection event* (CE). During a CE, data packets are exchanged between master and slave. After the CE, master and slave go to sleep mode to save energy. Exactly after the CI time period is over, both devices wake up to continue with a new CE (Figure C.1). The control and timing unit of this connection principle is located in the Link Layer. It is clocked by a low power low frequency clock domain. To ensure reliable functionality of this precisely timed schedule, no interruption by the application program is allowed. Therefore the Link Layer is decoupled from the Application Layer, which also ensures flexibility for application development.[1]

[1] Paragraph has already been published in [39]

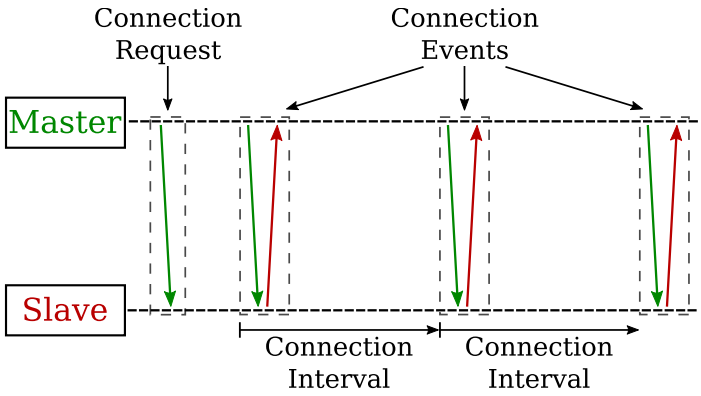

Figure C.1: Consecutive connection events, between which both devices sleep.[2]

In contrast to most other wireless communication technologies, BLE is optimized for low power operation and thus comes with a significantly lower data throughput. The design parameters that affect tradeoff between power demand and data throughput are the CI, the packet size and the number of packet per connection event as shown in Figure C.2.

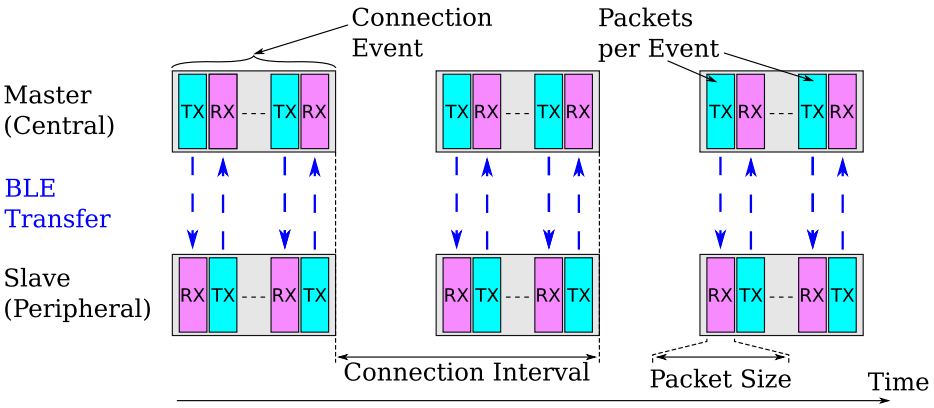

Figure C.2: Connection parameters of BLE that determine the throughput.[3]

Therefore, lowering the CI, increasing the packet size and increasing the number of packets per CE result in an increased data throughput but in return result in a higher energy demand.

[2]Graphic has already been published in [39]
[3]Graphic has already been published in [53]

Appendix D

Model Predictive Control on Field Programmable Gate Arrays

D.1 Basics of Model Predictive Control[1]

The basic idea of *Model Predictive Control* (MPC) is to predict the system behavior in the future using a system model

$$\boldsymbol{x}(k+1) = \boldsymbol{f}(\boldsymbol{x}(k), \boldsymbol{u}(k), k) \tag{D.1}$$

where k is the current time step, $\boldsymbol{x}(k)$ is the current system state, $\boldsymbol{u}(k)$ is the current system input and $\boldsymbol{x}(k+1)$ is the system state in the next consecutive time step. By solving the system model the future state sequence

$$\boldsymbol{X}(k) = \begin{bmatrix} \boldsymbol{x}(k+1) \\ \boldsymbol{x}(k+2) \\ \vdots \\ \boldsymbol{x}(k+N) \end{bmatrix} \tag{D.2}$$

is predicted within a specific prediction horizon N.

In order to obtain the optimal desired system behavior, an optimization problem is formulated whose solution points out the optimal system input sequence

$$\boldsymbol{U}(k) = \begin{bmatrix} \boldsymbol{u}(k) \\ \boldsymbol{u}(k+1) \\ \vdots \\ \boldsymbol{u}(k+N-1) \end{bmatrix} \tag{D.3}$$

[1] This section is based on the related lecture by Daniel Görges [296]

for the next consecutive time step. Therefore, a cost function

$$V_N(\boldsymbol{x}(k), \boldsymbol{U}(k)) = \sum_{i=0}^{N-1} \ell(\boldsymbol{x}(k+i), \boldsymbol{u}(k+i), k+i) \tag{D.4}$$

is formulated that is to be minimized

$$\begin{aligned} &\min_{\boldsymbol{U}(k)} V_N(\boldsymbol{x}(k), \boldsymbol{U}(k)) \\ &subject\ to \qquad\qquad \boldsymbol{x}(k+i+1) = \boldsymbol{f}\left(\boldsymbol{x}(k+i), \boldsymbol{u}(k+i), k+i\right) \end{aligned} \tag{D.5}$$

The optimization problem allows to consider specific state constrains $\boldsymbol{x}(k+i) \in \mathbb{X}(k+i)$ and input constrains $\boldsymbol{u}(k+i) \in \mathbb{U}(k+i)$.

Solving the optimization problem then points to the optimal input sequence

$$\boldsymbol{U}^*(k) = \begin{bmatrix} \boldsymbol{u}^*(k) \\ \boldsymbol{u}^*(k+1) \\ \vdots \\ \boldsymbol{u}^*(k+N-1) \end{bmatrix} \tag{D.6}$$

of which only the first input $\boldsymbol{u}^*(k)$ is applied to the system in the current time step. This way, a continuous feedback is present that ensures robustness against external disturbances and uncertainties of the system model.

Figure D.1 outlines the operations for an explanatory control system.

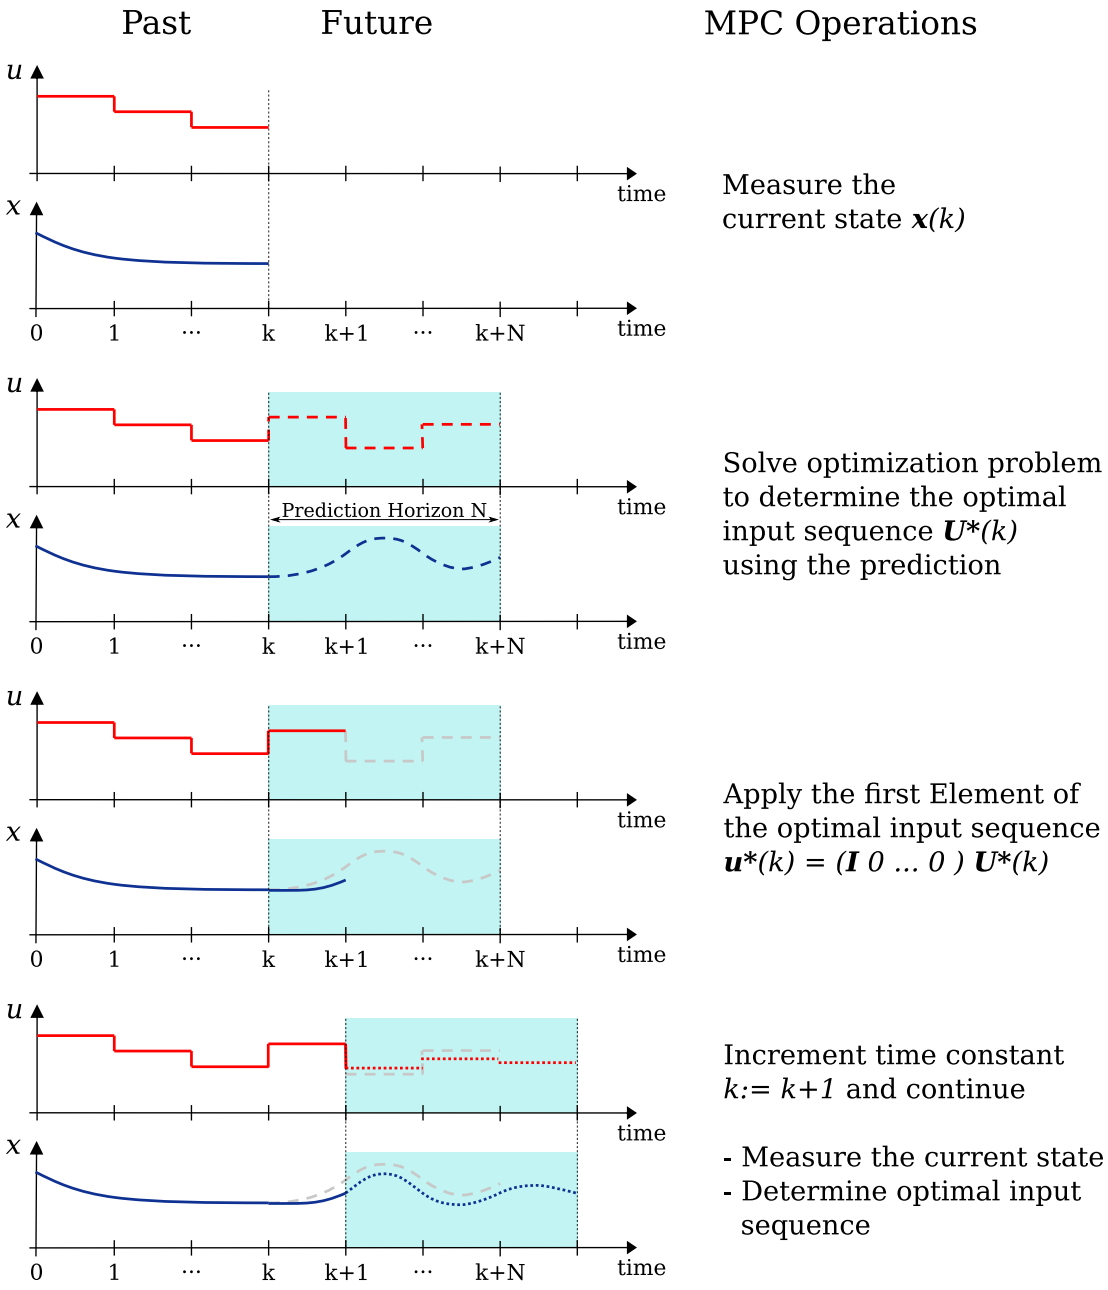

Figure D.1: Typical operations of an MPC system (graphic based on [296]).

D.2 Controller System Design

This section describes the components that have not been contributed by the author of this work but by its colleagues. These works have been a prerequisite to develop the hardware design strategy for the embedded MPC system.

D.2.1 MPC Controller Setup[2]

For this work, the control of linear discrete-time systems of the form

$$\boldsymbol{x}_{t+1} = \boldsymbol{A}\boldsymbol{x}_t + \boldsymbol{B}\boldsymbol{u}_t + \boldsymbol{d}_t \tag{D.7}$$

is considered, where $\boldsymbol{x} \in \mathbb{R}^n$ and $\boldsymbol{u} \in \mathbb{R}^m$ are the state and input of the system. $\boldsymbol{d} \in \mathbb{R}^n$ is a known time-varying exogenous input. The matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are the system and the input matrix, respectively. The state and input are subjects to the constraints

$$\boldsymbol{x}_{\min} \leq \boldsymbol{x}_t \leq \boldsymbol{x}_{\max}, \quad \boldsymbol{u}_{\min} \leq \boldsymbol{u}_t \leq \boldsymbol{u}_{\max}. \tag{D.8}$$

The embedded implementation of the following MPC control problem is studied in this work

$$\begin{aligned}
&\underset{\boldsymbol{X}_t, \boldsymbol{U}_t}{\text{minimize}} && V_{N,t}\left(\boldsymbol{X}_t, \boldsymbol{U}_t\right) = \sum_{r=t}^{t+N-1} l_r\left(\boldsymbol{x}_r, \boldsymbol{u}_r\right) && \text{(D.9a)}\\
&\text{subject to} && \boldsymbol{x}_t = \boldsymbol{x}_t^{\mathrm{m}} && \text{(D.9b)}\\
& && \text{for all } r \in \mathcal{I}_{t:t+N-2}: && \\
& && \boldsymbol{x}_{r+1} = \boldsymbol{A}\boldsymbol{x}_r + \boldsymbol{B}\boldsymbol{u}_r + \boldsymbol{d}_r && \text{(D.9c)}\\
& && \text{for all } r \in \mathcal{I}_{t:t+N-1}: && \\
& && \boldsymbol{x}_{\min} \leq \boldsymbol{x}_r \leq \boldsymbol{x}_{\max} && \text{(D.9d)}\\
& && \boldsymbol{u}_{\min} \leq \boldsymbol{u}_r \leq \boldsymbol{u}_{\max} && \text{(D.9e)}
\end{aligned}$$

where N is the prediction horizon, $\boldsymbol{X}_t = \left[\boldsymbol{x}_t^{\mathrm{T}}, \ldots, \boldsymbol{x}_{t+N-1}^{\mathrm{T}}\right]^{\mathrm{T}}$ and $\boldsymbol{U}_t = \left[\boldsymbol{u}_t^{\mathrm{T}}, \ldots, \boldsymbol{u}_{t+N-1}^{\mathrm{T}}\right]^{\mathrm{T}}$. In (D.9c) the plant model (D.7) is used to predict the future behavior of the plant starting from the measurement $\boldsymbol{x}_t^{\mathrm{m}}$. With (D.9d) and (D.9e) the constraints (D.8) are considered. In the optimization problem the cost function $V_{N,t}$ is minimized. It is the sum of the time-varying stage cost

$$l_t\left(\boldsymbol{x}_t, \boldsymbol{u}_t\right) = \boldsymbol{x}_t^{\mathrm{T}}\boldsymbol{Q}_t\boldsymbol{x}_t + \boldsymbol{q}_t^{\mathrm{T}}\boldsymbol{x}_t + \boldsymbol{u}_t^{\mathrm{T}}\boldsymbol{R}_t\boldsymbol{u}_t + \boldsymbol{r}_t^{\mathrm{T}}\boldsymbol{u}_t$$

over the prediction horizon. $\boldsymbol{Q}_t$ and $\boldsymbol{R}_t$ are positive definite matrices for all $t \in \mathbb{N}$.

[2]The controller setup has already been published in [48]

The MPC problem (D.9) must be solved online at every time instant t. It can be recast as a strictly convex quadratic program (QP) of the form

$$\begin{aligned} \underset{\boldsymbol{z}}{\text{minimize}} \quad & \frac{1}{2}\boldsymbol{z}^{\mathrm{T}}\boldsymbol{H_t}\boldsymbol{z} + \boldsymbol{f_t}^{\mathrm{T}}\boldsymbol{z} + c && \text{(D.10a)} \\ \text{subject to} \quad & \boldsymbol{A}_{\text{eq}}\boldsymbol{z} = \boldsymbol{B}_{\text{eq}} && \text{(D.10b)} \\ & \boldsymbol{A}_{\text{in}}\boldsymbol{z} \leq \boldsymbol{B}_{\text{in}}, && \text{(D.10c)} \end{aligned}$$

where $\boldsymbol{H_t}$ is a positive definite matrix. $\boldsymbol{z} = \left[\boldsymbol{x}_t^{\mathrm{T}}, \boldsymbol{u}_t^{\mathrm{T}}, \ldots, \boldsymbol{x}_{t+N-1}^{\mathrm{T}}, \boldsymbol{u}_{t+N-1}^{\mathrm{T}}\right]^{\mathrm{T}}$ is the optimization vector and

$$\boldsymbol{H_t} = \begin{bmatrix} \boldsymbol{Q}_t & & & \\ & \boldsymbol{R}_t & & \\ & & \ddots & \\ & & & \boldsymbol{R}_{t+N-1} \end{bmatrix}, \boldsymbol{f_t} = \begin{bmatrix} \boldsymbol{q}_t \\ \boldsymbol{r}_t \\ \vdots \\ \boldsymbol{r}_{t+N-1} \end{bmatrix}$$

$$\boldsymbol{A}_{\text{eq}} = \begin{bmatrix} \boldsymbol{I} & & & & & \\ \boldsymbol{A} & \boldsymbol{B} & -\boldsymbol{I} & & & \\ & & \ddots & \ddots & \ddots & \\ & & & \boldsymbol{A} & \boldsymbol{B} & -\boldsymbol{I} \end{bmatrix}, \boldsymbol{B}_{\text{eq}} = \begin{bmatrix} \boldsymbol{x}_t^m \\ -\boldsymbol{d}_t \\ \vdots \\ -\boldsymbol{d}_{t+N-2} \end{bmatrix}$$

$$\boldsymbol{A}_{\text{in}} = \begin{bmatrix} \boldsymbol{I} & & & \\ -\boldsymbol{I} & & & \\ & \ddots & & \\ & \ddots & & \\ & & \boldsymbol{I} & \\ & & -\boldsymbol{I} & \end{bmatrix}, \boldsymbol{B}_{\text{in}} = \begin{bmatrix} \boldsymbol{z}_{\max} \\ \boldsymbol{z}_{\min} \\ \vdots \\ \vdots \\ \boldsymbol{z}_{\max} \\ \boldsymbol{z}_{\min} \end{bmatrix}$$

$$\boldsymbol{z}_{\max} = \left[\boldsymbol{x}_{\max}^{\mathrm{T}}, \boldsymbol{u}_{\max}^{\mathrm{T}}\right]^{\mathrm{T}}, \boldsymbol{z}_{\min} = -\left[\boldsymbol{x}_{\min}^{\mathrm{T}}, \boldsymbol{u}_{\min}^{\mathrm{T}}\right]^{\mathrm{T}}.$$

There are many algorithms in the literature for solving the QP (D.10), see [297] for an overview. In this paper Algorithm D.1, which consists of a fast dual ascent method, similar to the one in [298], is applied.

k is the iteration index, $k_{\max}$ maximum number of iterations. $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ are the Karush-Kuhn-Tucker (KKT) multipliers of the equality and the inequality constraints of the dual problem, respectively. $\nu_k = \frac{k-1}{k+2}$ is an iteration-varying scalar. $\boldsymbol{L}_\mu$ and $\boldsymbol{L}_\lambda$ are scaling matrices that precondition the problem and can be determined offline. To preserve the structure in the algorithm and allow a computationally efficient implementation, $\boldsymbol{L}_\mu$ is chosen as

$$\boldsymbol{L}_\mu = \boldsymbol{I}_{(N-1)\times(N-1)} \otimes \tilde{\boldsymbol{L}}_\mu$$

where $\tilde{\boldsymbol{L}}_\mu$ is a diagonal matrix and $\otimes$ denotes the Kronecker product. Similarly, $\boldsymbol{L}_\lambda$ is chosen as a block-diagonal matrix $\boldsymbol{L}_\lambda = \boldsymbol{I}_{(N-1)\times(N-1)} \otimes \tilde{\boldsymbol{L}}_\lambda$.

Algorithm D.1: *Fast Dual Ascent* (FDA) algorithm.

Require: $\boldsymbol{z}^{-1}, \boldsymbol{\lambda}^{0}, \boldsymbol{\lambda}^{-1}, \boldsymbol{\mu}^{0}, \boldsymbol{\mu}^{-1}$

for $k = 0, \ldots, k_{\max}$

$$\begin{aligned}
\boldsymbol{z}^{k} &:= -\boldsymbol{H}^{-1}\left(\begin{bmatrix}\boldsymbol{A}_{\text{eq}}^{\text{T}} & \boldsymbol{A}_{\text{in}}^{\text{T}}\end{bmatrix}\begin{bmatrix}\boldsymbol{\lambda}^{k} \\ \boldsymbol{\mu}^{k}\end{bmatrix} + \boldsymbol{f}\right) \\
\bar{\boldsymbol{z}}^{k} &:= \boldsymbol{z}^{k} + \nu_k\left(\boldsymbol{z}^{k} - \boldsymbol{z}^{k-1}\right) \\
\boldsymbol{\lambda}^{k+1} &:= \boldsymbol{\lambda}^{k} + \nu_k\left(\boldsymbol{\lambda}^{k} - \boldsymbol{\lambda}^{k-1}\right) + \boldsymbol{L}_{\lambda}^{-1}\left(\boldsymbol{A}_{\text{eq}}\bar{\boldsymbol{z}}^{k} - \boldsymbol{B}_{\text{eq}}\right) \\
\boldsymbol{\mu}^{k+1} &:= \max\{\boldsymbol{0}, \boldsymbol{\mu}^{k} + \nu_k\left(\boldsymbol{\mu}^{k} - \boldsymbol{\mu}^{k-1}\right) + \boldsymbol{L}_{\mu}^{-1}\left(\boldsymbol{A}_{\text{in}}\bar{\boldsymbol{z}}^{k} - \boldsymbol{B}_{\text{in}}\right)\}
\end{aligned}$$

end for

Although the problems of the form (D.9) is quite general and thus addresses many applications, the aim of this paper is the efficient implementation for controllers in infrastructure systems. Those systems can be controlled with a hierarchical scheme consisting of an upper level and lower level controller [299]. The upper level controller optimizes the economically behavior of the overall system over a long horizon (up to one day) using a large sampling time (i.e. several minutes or hours), while the lower level controller undertakes short-term control actions within a shorter time horizon. It is assumed that problem (D.9) represents an upper level control problem. Applications can be found in the control of power networks [86] or WDNs [258]. Here the exogenous input is usually a consumer demand that needs to be satisfied. The time-varying cost function stems from a time-varying price for power over the day.

For this work, an existing Matlab reference design of the discussed MPC controller was provided by Dr.-Ing. Felix Berkel.

D.2.2 WDN Benchmark

A *water distribution network* (WDN) system was employed as benchmark, which has been calibrated with real measurement data from a WDN provider. A detailed description can be found in [258]. This benchmark system consists of three water tanks, two pumps and two water sources, i.e., the system has three states and four inputs.

In a WDN, pumps transport water from the sources, such as springs or deep-wells, to the consumer. There are two locations of water consumption and one free source in the benchmark, which represents a river flowing into the system. The maximum error of the system was tuned in simulation to provide a compromise between closed-loop control performance and computational effort. It was determined, that the variance of each element of the optimized output needs to be lower than 10^{-2} in order to not affect the control performance.[3]

[3] Paragraph has already been published in [48]

D.3 Field Programmable Gate Arrays

The metrics that determine the performance of *Field Programmable Gate Array* (FPGA) designs are throughput, clock frequency and latency. The throughput is determined by the amount of input data that can be processed by the hardware. And the latency in clock cycles describes the number of clock cycles that a hardware requires to process this input until the output is affected by it. Therefore, the latency and the clock frequency determine the computational speed of the system, whereas combined with the throughput the overall performance is determined.

Parallelization techniques are used to increase the performance, either by creating multiple computation instances in parallel, which increases the throughput due to a wider input bit width or by pipelining, which increases the throughput by a more frequent input data acquisition.

Pipelining is a principle to avoid data dependencies between consecutive single computation units of a sequential computation stage. By inserting registers in between the single computation units, the intermediate results are decoupled from each other. This way, all computation units are executing concurrently in each clock cycle, shift their intermediate result to the next stage at the concurrent clock cycle and fetch the result from its predecessor unit. Figure D.2 shows the difference between an implementation without and one with pipelining.

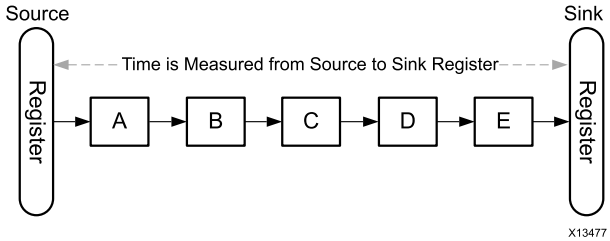

FPGA Implementation *without* Pipelining

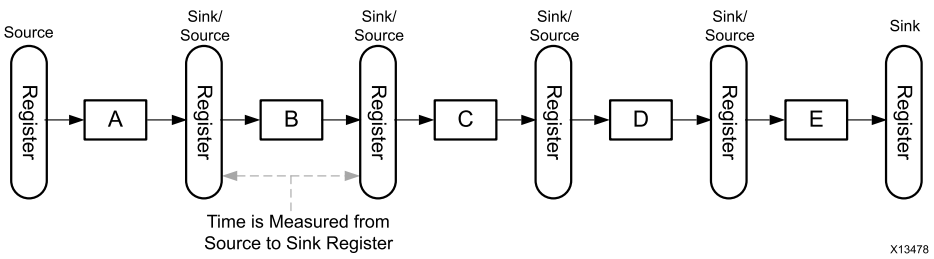

FPGA Implementation *with* Pipelining

Figure D.2: Pipelining of a computation stage that consist of multiple sequential computation units (A...E) (source: Xilinx [236]).

Because of the additional registers, pipelining increases the latency in general. But, in return, the data path lengths and the respective signal propagation time is reduced and

therefore the clock frequency can be increased. Furthermore, it increases the throughput as new input data can be fetched faster one after another. When applied to iterative computations in loops, pipelining can influence the number of clock cycles between consecutive loop iterations, also known as *initialization interval* (II). There are further parallelization techniques for loops, such as unrolling and flattening. However, these come at the expense of even more additional hardware resources for additionally required computation units and their usage is contingent on data dependencies between the loop iteration steps. Figure D.3 shows the relation between different parallelization techniques, latency and required hardware resources for a matrix vector multiplication.

$$\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = \begin{pmatrix} a_{11} * b_1 + a_{12} * b_2 \\ a_{21} * b_1 + a_{22} * b_2 \end{pmatrix}$$

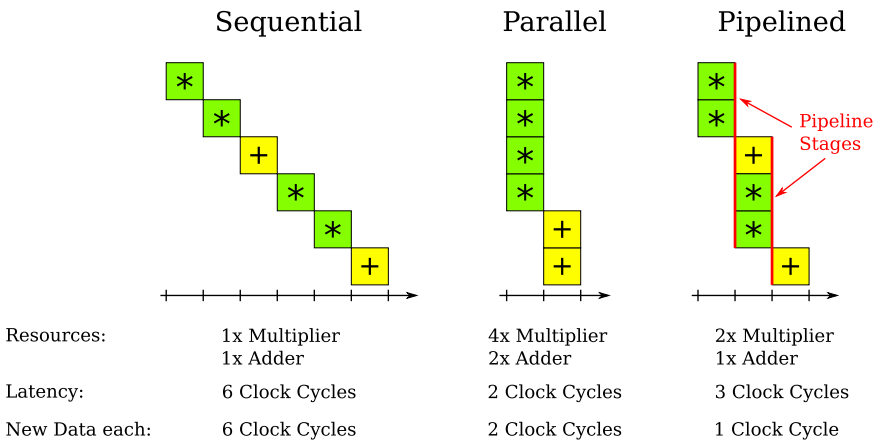

Figure D.3: Comparison of parallelization techniques on FPGAs and the respective latency and required hardware resources.

The design tradeoffs like the one between hardware resource requirements and execution time offer a high degree of flexibility to the hardware designer. The choice of the number format furthermore enables a straight tradeoff between required hardware resources, latency and computational precision. More specifically: The bit length of the chosen number format directly determines the demand of memory cells as well as the complexity and bit width of the calculation units, which is reflected in the demand of arithmetic and logic cells. Furthermore, the latency is affected as computations for number formats with lower bit length are usually faster than for those with larger bit lengths.

Besides the bit length, the type of number representation has a significant effect on hardware resource demand and latency. For instance, fixed-point number formats require less complex calculation units than floating-point number formats and therefore usually show a lower latency and lower hardware resources for the same bit length. However, fixed-point formats in return cover bot a significantly lower range and lower precision, so that the precision of the computation outputs is usually lower than those of floating-point format-based designs.

Acronyms

3GPP	3rd Generation Partnership Project. 81, 82
6LoWPAN	IPv6 over Low Power Wireless Personal Area Network. 84, 85, 87, 89, 90
AC	alternating current. 9, 22, 23, 130, 131
ADC	analog to digital converter. 46
AI	artificial intelligence. 114, 128
ALU	Arithmetic Logic Unit. 96
API	application programming interface. 26
ASIC	application-specific integrated circuit. 1
BAN	Body Area Network. 63, 64, 73, 74, 173
BLE	Bluetooth Low Energy. 11, 15, 59, 63, 64, 66, 69, 70, 72, 73, 77, 85, 87, 88, 89, 90, 112, 119, 121, 127, 137, 138, 173, 175
BRAM	Block RAM. 96, 98
CAGR	Compound Annual Growth Rate. 4
CE	connection event. 137, 138
CI	connection interval. 137, 138
CIoT	Cellular IoT. 81, 82
CLB	Configurable Logic Block. 96
CP	Checkpointing. 25, 26, 27, 28, 29, 30, 32, 37, 43, 49, 53
CPU	Central Processing Unit. 1, 2, 12, 27, 28, 76, 87, 90, 95, 97, 98, 113, 119
CSU	Cold-Start Unit. 130
DC	direct current. 22, 52, 53, 54, 55, 58, 132, 133
DDR	Double Data Rate. 16, 123
DH	Diffie-Hellman. 83
DHE	Diffie-Hellman ephemeral. 83, 91
DIMM	Dual Inline Memory Module. 16, 123
DRAM	Dynamic Random Access Memory. 16, 123
DSA	Digital Signature Algorithm. 83, 91
DSP	Digital Signal Processor. 1, 96

ECC	Elliptic Curve Cryptography. 84, 90, 91
EEPROM	electrically erasable programmable read-only memory. 25, 88
FDA	Fast Dual Ascent. 144, 179
FF	Flip Flop. 96, 97
F-PLD	frequency-based power loss detection. 43, 44, 172
FPGA	Field Programmable Gate Array. 1, 12, 15, 93, 94, 95, 96, 97, 98, 99, 100, 101, 106, 108, 113, 119, 120, 122, 145, 146, 174, 175
FRAM	ferro-electric RAM. 25, 26, 31, 32, 134, 135
gmPSD	guaranteed minimal power supply duration. 42, 43, 172
GPU	Graphics Processing Unit. 1, 16, 95, 98, 123
HDL	Hardware Description Language. 97, 98, 101
HLS	High Level Synthesis. 97, 98, 101, 107
IC	integrated circuit. 1, 3, 51, 54, 55, 127
IC	integrated circuit. 53, 54
II	initialization interval. 146
IIoT	Industrial Internet of Things. 4, 83
IMU	Inertial Measurement Unit. 14, 61, 62, 65, 66, 67, 75, 121, 173
IoT	Internet of Things. 3, 4, 7, 10, 11, 15, 24, 32, 45, 46, 48, 50, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 112, 113, 118, 119, 122, 127, 128, 131, 173, 174
IP	intellectual property. 3
IP	inflection point. 38, 40, 41, 42, 43, 172
IP	Internet Protocol. 80, 83, 84, 85, 86, 87, 89, 90
IPv6	Internet Protocol v6. 84, 85, 87, 90
IRQ	interrupt request. 40, 70, 71, 173
LAN	Local Area Network. 82, 84
LCD	liquid crystal display. 129
LIB	Lithium-Ion battery. 5, 7
LoRa-WAN	Long Range WAN. 82, 84
LTE	Long Term Evolution. 81, 82
LTE-M	LTE-MTC (Machine Type Communication). 82
LUT	look-up table. 42, 96, 97, 105

MCU	Microcontroller Unit. 1, 3, 4, 19, 21, 23, 24, 25, 26, 27, 28, 30, 31, 32, 33, 34, 35, 36, 40, 41, 42, 43, 51, 52, 53, 54, 55, 56, 65, 66, 80, 127, 129, 130, 134, 135, 171
MEMS	micro electromechanical system. 14, 61, 75, 121, 128, 129, 174
MIMO	Multiple Input Multiple Output. 82
ML	machine learning. 114, 128
MPC	Model Predictive Control. 11, 12, 15, 93, 94, 95, 97, 98, 99, 100, 101, 102, 105, 106, 107, 108, 109, 113, 119, 122, 139, 141, 142, 144, 175
MPP	maximum power point. 33, 34, 57, 58, 133
MPPT	maximum power point tracking. 21, 33, 34, 35, 132, 133
MPSoC	multiprocessor SoC. 109
MQTT	Message Queuing Telemetry Transport. 89, 90, 91, 92, 173, 174
MRAM	Magnetoresistive RAM. 25, 114, 135, 136
NB-IoT	Narrowband-IoT. 82
NEMS	Nanoelectromechanical system. 128
NFC	Near Field Communication. 133
NoC	Network on Chip. 1
NVM	nonvolatile memory. 21, 25, 26, 27, 28, 29, 30, 32, 37, 134, 135, 136
NVP	Nonvolatile Processor. 26, 27, 30
PAN	Personal Area Network. 73, 82, 84, 89, 127
PC	Personal Computer. 3, 80
PCB	printed circuit board. 67, 88
PdWU	power-down and wake-up unit. 87, 88, 92
PFS	Perfect Forward Secrecy. 83, 85, 90, 91, 92
PKI	Public Key Infrastructure. 83
PL	programmable logic. 109
PLD	power loss detection. 27, 28, 30, 31, 32, 42, 44
PS	processing system. 100, 109
PSD	power supply duration. 40, 41, 42, 43
RAM	random access memory. 25, 27, 28, 54, 56
RF	radio frequency. 16, 24, 75, 123, 127, 133
RFID	radio frequency identification. 24, 133
RRAM	Resistive RAM. 26, 135
RSA	Rivest–Shamir–Adleman. 84
RTC	Real-Time Clock. 87, 88
SDC	Structural Data Compression. 106, 108, 113, 174
SDK	Software Development Kit. 87, 90

SIM	Subscriber Identity Module. 82
SoC	System on Chip. 1, 32, 46, 47, 48, 49, 70, 71, 74, 75, 76, 87, 88, 89, 90, 127, 173
SRAM	static random access memory. 26
STT	Spin Transfer Torque. 114
TCP	Transmission Control Protocol. 80, 83, 89
TDMA	time division multiple access. 74
TEG	thermoelectric generator. 33, 45, 46, 47, 49, 50, 131, 132, 172
TLS	Transport Layer Security. 11, 15, 80, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 112, 119, 122, 173, 174
UDP	User Datagram Protocol. 89
USB	Universal Serial Bus. 66, 73
UVC	Undervoltage Condition. 52, 53, 54, 55, 56
VC	voltage converter. 41, 42
VM	voltage monitor. 42
V-PLD	voltage-based power loss detection. 41, 42, 43, 44, 172
WAN	Wide Area Network. 81, 82, 89
WDN	water distribution network. 100, 101, 144
WiFi	wireless fidelity (WLAN). 84, 127
WLAN	Wireless LAN. 84
WSN	wireless sensor network. 11, 59, 60, 64, 77, 112, 127
ZIPS	zetta instructions per second. 4, 5, 171

Bibliography

[1] G. E. Moore. Cramming More Components Onto Integrated Circuits. *Electronics*, page 114–117, April 1965.

[2] S. R. C. SRC. Decadal Plan for Semiconductors. Technical report, Semiconductor Research Corporation (SRC), 2021.

[3] MIT Climate & Sustainability Consortium. How can we reduce the carbon footprint of global computing? `https://climate.mit.edu/posts/how-can-we-reduce-carbon-footprint-global-computing (26.04.2023)`, April 2022.

[4] S. Ahmed, B. Islam, K. S. Yildirim, M. Zimmerling, P. Pawełczak, M. H. Alizai, B. Lucia, L. Mottola, J. Sorber, and J. Hester. The Internet of Batteryless Things. *Commun. ACM*, 67(3):64–73, feb 2024.

[5] J. Eitzinger. Evaluation of Intel Xeon Phi "knights landing": Initial impressions and benchmarking results. In *Intel MIC Programming Workshop Proceedings*, 2017.

[6] J. Hennessy and D. Patterson. *Computer Architecture: A Quantitative Approach*. ISSN. Elsevier Science, 2017.

[7] Hannah Ritchie, Max Roser. Moore's Law: The number of transistors on microchips doubles every two years. `https://ourworldindata.org/uploads/2020/11/Transistor-Count-over-time.png (01.04.2022)`, November 2020.

[8] anbofer. EVOLUTION OF CPU (HISTORY). `https://www.timetoast.com/timelines/evolution-of-cpu-history (01.04.2022)`, 2017.

[9] R. Dennard, V. Rideout, E. Bassous, and A. LeBlanc. Design of ion-implanted MOSFET's with very small physical dimensions. *IEEE Journal of Solid-State Circuits*, 9(5):256–268, 1974.

[10] Statista. Internet of Things (IoT) - statistics & facts. `https://www.statista.com/topics/2637/internet-of-things/#topicOverview (26.04.2023)`, 2022.

[11] Mikatech. Microcontrollers: A Designer's Overview. `https://www.break-ic.com/topics/break-microcontroller.asp (21.01.2022)`, 2013.

[12] Total Phase, Inc. How Embedded Systems Impact Your Everyday Life. `https://www.totalphase.com/blog/2018/01/embedded-systems-impact-everyday-life/ (21.01.2022)`, January 2018.

[13] Microchip. What is a Microcontroller? `http://ww1.microchip.com/downloads/en/Market_Communication/mcu.pdf (21.01.2022)`.

[14] Research and Markets. Microcontroller Market: Global Industry Trends, Share, Size, Growth, Opportunity and Forecast 2021-2026. `https://www.`

researchandmarkets.com/reports/5331086/microcontroller-market-global-industry-trends (21.01.2022), May 2021.

[15] Fact.MR. Microcontroller Market Share and Trends Analysis [...] Forecast, 2021-2031. https://www.factmr.com/report/4556/microcontroller-market (21.01.2022), October 2021.

[16] Grand View Research. Microcontroller Market Size, Share and Trends Analysis Report [...] Forecasts, 2021 - 2028. https://www.grandviewresearch.com/industry-analysis/microcontroller-market (21.01.2022), March 2021.

[17] Grand View Research. Microcontroller Market Size, Share and Trends Analysis Report [...] Forecasts, 2023 - 2030. https://www.grandviewresearch.com/industry-analysis/microcontroller-market (16.03.2023), March 2022.

[18] Global Market Insights Inc. Microcontroller Market Size [...] Competitive Market Share and Forecast, 2021 – 2027. https://www.gminsights.com/industry-analysis/microcontroller-market (21.01.2022), 2020.

[19] Market Data Forecast. Microcontroller Market Analysis [...] Growth, Trends, Forecast — 2021 – 2026. https://www.marketdataforecast.com/market-reports/global-microcontroller-market (21.01.2022), April 2021.

[20] Data Bridge Market Research. Global Microcontroller Market – Industry Trends and Forecast to 2028. https://www.databridgemarketresearch.com/reports/global-microcontroller-market (21.01.2022), June 2021.

[21] Research Reports World. Global Microcontrollers (MCU) Market — Growing at CAGR 5.1% — Expected to Reach USD 35950 Million — Forecast Period 2022-2027. https://www.rfdtv.com/story/45269218/global-microcontrollers-mcu-market-growing-at-cagr-51-expected-to-reach-usd-35950-million-forecast-period-2022-2027 (21.01.2022), November 2021.

[22] Precedence Research. Microcontroller (MCU) Market [...] - Global Market Size, Trends Analysis, Segment Forecasts, Regional Outlook 2020 - 2027. https://www.precedenceresearch.com/microcontroller-mcu-market (21.01.2022), 2021.

[23] GlobeNewswire. Microcontroller Market size worth 786.90 Million Dollar, Globally, by 2028 at 14.20% CAGR: Verified Market Research. https://www.globenewswire.com/news-release/2021/11/11/2332750/0/en/Microcontroller-Market-size-worth-786-90-Million-Globally-by-2028-at-14-20-CAGR-Verified-Market-Research.html (21.01.2022), November 2021.

[24] Statista. Number of Internet of Things (IoT) connected devices worldwide from 2019 to 2021, with forecasts from 2022 to 2030. https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/ (26.04.2023), 2022.

[25] ScienceAlert. Computers Will Require More Energy Than The World Generates by 2040. https://www.sciencealert.com/computers-will-require-more-energy-than-the-world-generates-by-2040 (26.04.2023), July 2016.

[26] Statista. Lithium-ion batteries - statistics and facts. https://www.statista.com/topics/2049/lithium-ion-battery-industry/#dossierKeyfigures (20.12.2021), 2021.

[27] Forever EV. Global Lithium Battery Demand and Capacity Forecast by Sector 2020-2030. https://www.foreverev.com/2021/08/global-lithium-battery-demand-and-capacity-forecast-by-sector-2020-2030/ (20.12.2021), 2021.

[28] Statista. High Demand for Lithium-Ion Batteries. https://www.statista.com/chart/23808/lithium-ion-battery-demand/ (20.12.2021), 2020.

[29] G. Zubi, R. Dufo-López, M. Carvalho, and G. Pasaoglu. The lithium-ion battery: State of the art and future perspectives. *Renewable and Sustainable Energy Reviews*, 89:292–308, 2018.

[30] E. Fan, L. Li, Z. Wang, J. Lin, Y. Huang, Y. Yao, R. Chen, and F. Wu. Sustainable Recycling Technology for Li-Ion Batteries and Beyond: Challenges and Future Prospects. *Chemical Reviews*, 120(14):7020–7063, 2020. PMID: 31990183.

[31] J. Cheng, B. Zhou, K. Kunze, C. C. Rheinländer, S. Wille, N. Wehn, J. Weppner, and P. Lukowicz. Activity Recognition and Nutrition Monitoring in Every Day Situations with a Textile Capacitive Neckband. In *Proceedings of the 2013 ACM Conference on Pervasive and Ubiquitous Computing Adjunct Publication*, UbiComp '13 Adjunct, page 155–158, New York, NY, USA, 2013. Association for Computing Machinery.

[32] T. Jaitner, M. Schmidt, K. Nolte, C. Rheinländer, S. Wille, and N. Wehn. A Wearable Inertial Sensor Unit for Jump Diagnosis in Multiple Athletes. In *icSPORTS 2014 - Proceedings of the 2nd International Congress on Sports Sciences Research and Technology Support*, 10 2014.

[33] M. Schmidt, T. Jaitner, C. Rheinländer, S. Wille, and N. Wehn. Sprungdiagnostik unter Feldbedingungen mittels Intertialsensoren. In *Proceedings of the 10. Symposium der dvs-Sektion Sportinformatik*, 2014.

[34] T. Jaitner, C. Rheinländer, M. Schmidt, S. Wille, and N. Wehn. Eine flexible WSN-Plattform zur Leistungsdiagnostik und Trainingssteuerung im Leistungs- und Hochleistungssport. *16. Frühjahrsschule "Informations- und Kommunikationstechnologien in der angewandten Trainingswissenschaft"*, 2014.

[35] T. Jaitner, M. Schmidt, K. Nolte, C. Rheinländer, S. Wille, and N. Wehn. Vertical jump diagnosis for multiple athletes using a wearable inertial sensor unit. *Sports Technology*, 8(1-2):51–57, 2015.

[36] M. Schmidt, T. Jaitner, J.-E. Jürjens, K. Nolte, C. Rheinländer, S. Wille, and N. Wehn. Entwicklung und Evaluation eines mobilen Messsystems zum Online-Monitoring von Leistungsparametern hochdynamischer Sprint- und Sprungbewegungen (SpoSeNs). Technical report, BISp-Jahrbuch Forschungsförderung 2015/16, 2015.

[37] R. Xiong, S. van Waasen, J. Schelten, M. Schloesser, C. Rheinländer, and N. Wehn. Development of a quasi time stretch technology for indoor positioning system based on pulse modulated ultra high frequency radio. In *2015 IEEE SENSORS*, pages 1–4, 2015.

[38] M. Jung, C. Rheinländer, C. Weis, and N. Wehn. Reverse Engineering of DRAMs: Row Hammer with Crosshair. In *International Symposium on Memory Systems (MEMSYS 2016)*, 2016.

[39] C. C. Rheinländer and N. Wehn. Precise synchronization time stamp generation for Bluetooth low energy. In *2016 IEEE SENSORS*, pages 1–3, 2016.

[40] M. Schmidt, C. C. Rheinländer, S. Wille, N. Wehn, and T. Jaitner. IMU-Based Determination of Fatigue during Long Sprint. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, UbiComp '16, page 899–903, New York, NY, USA, 2016. Association for Computing Machinery.

[41] M. Schmidt, C. Rheinländer, K. F. Nolte, S. Wille, N. Wehn, and T. Jaitner. IMU- based Determination of Stance Duration During Sprinting. *Procedia Engineering*, 147:747 – 752, 2016. The Engineering of SPORT 11.

[42] R. Xiong, S. van Waasen, C. Rheinlnder, and N. Wehn. Development of a Novel Indoor Positioning System With mm-Range Precision Based on RF Sensors Network. *IEEE Sensors Letters*, 1(5):1–4, 2017.

[43] M. Jung, D. Mathew, C. Rheinländer, C. Weis, and N. Wehn. A Platform to Analyze DDR3 DRAM's Power and Retention Time. *IEEE Design & Test*, 34(4):52–59, Aug 2017.

[44] D. M. Mathew, M. Schultheis, C. C. Rheinländer, C. Sudarshan, M. Jung, C. Weis, and N. Wehn. An Analysis on Retention Error Behavior and Power Consumption of Recent DDR4 DRAMs. In *IEEE Conference Design, Automation and Test in Europe (DATE)*, 2018.

[45] M. Jung, D. M. Mathew, C. C. Rheinländer, C. Weis, and N. Wehn. A Platform for Analyzing DDR3 and DDR4 DRAMs. In *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2018.

[46] M. Schmidt, S. Wille, C. Rheinländer, N. Wehn, and T. Jaitner. A Wearable Flexible Sensor Network Platform for the Analysis of Different Sport Movements. In T. Ahram and C. Falcão, editors, *Advances in Human Factors in Wearable Technologies and Game Design*, pages 3–14, Cham, 2018. Springer International Publishing.

[47] C. C. Rheinländer and N. Wehn. Adaptive Transient Computing for Power-Neutral Embedded Devices. In *2019 29th International Symposium on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, pages 63–68, 2019.

[48] C. C. Rheinländer, F. Berkel, M. Douglas, M. Schäfer, N. Wehn, and S. Liu. Structural Data Compression for Embedded Long Prediction Horizon Model Predictive Control on Resource-Constrained FPGA Platforms. In *2019 1st International Conference on Electrical, Control and Instrumentation Engineering (ICECIE)*, pages 1–7, 2019.

[49] M. D. da Silva Carvalho, F. Koark, C. Rheinländer, and N. Wehn. Real-Time Image Recognition System Based on an Embedded Heterogeneous Computer and Deep Convolutional Neural Networks for Deployment in Constrained Environments. In *WCX SAE World Congress Experience*. SAE International, apr 2019.

[50] S. Kapp, M. Thees, M. P. Strzys, F. Beil, J. Kuhn, O. Amiraslanov, H. Javaheri, P. Lukowicz, F. Lauer, C. Rheinländer, and N. Wehn. Augmenting Kirchhoff's laws: Using augmented reality and smartglasses to enhance conceptual electrical experiments for high school students. *The Physics Teacher*, 57(1):52–53, 2019.

[51] C. C. Rheinländer and N. Wehn. Harvester-aware transient computing: Utilizing the mechanical inertia of kinetic energy harvesters for a proactive frequency-based power loss detection. *Integration*, 75:122 – 130, 2020.

[52] J. Mades, G. Ebelt, B. Janjic, F. Lauer, C. Rheinländer, and N. Wehn. TLS-Level Security for Low Power Industrial IoT Network Infrastructures. In *2020 IEEE Conference Design, Automation and Test in Europe (DATE)*. IEEE, to appear in March 2020. preprint on webpage at `https://ems.eit.uni-kl.de/fileadmin/ems/user_upload/low_`

`power_iot_security_preprint.pdf` as well as `https://www.frederiklauer.de/lowPowerIotSecurityPreprint.pdf`.

[53] F. Lauer, C. C. Rheinländer, C. Kestel, and N. Wehn. Analysis and Optimization of TLS-based Security Mechanisms for Low Power IoT Systems. In *2020 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID)*, pages 775–780, 2020.

[54] M. Schöffel, F. Lauer, C. C. Rheinländer, and N. Wehn. On the Energy Costs of Post-Quantum KEMs in TLS-Based Low-Power Secure IoT. In *Proceedings of the International Conference on Internet-of-Things Design and Implementation*, IoTDI '21, page 158–168, New York, NY, USA, 2021. Association for Computing Machinery.

[55] M. Schöffel, F. Lauer, C. C. Rheinländer, and N. Wehn. Secure IoT in the Era of Quantum Computers—Where Are the Bottlenecks? *Sensors*, 22(7), 2022.

[56] S. Kapp, F. Lauer, F. Beil, C. C. Rheinländer, N. Wehn, and J. Kuhn. Smart Sensors for Augmented Electrical Experiments. *Sensors*, 22(1), 2022.

[57] S. Thielen, J. Seewig, N. Wehn, Y. Stiemcke, P. Sivasothy, C. Rheinländer, B. Sauer, and O. Koch. Sensorintegrierendes Maschinenelement Radialwellendichtring zur Überwachung von Betriebszustand und Schmierstoffqualität. In *Dresdner Maschinenelemente Kolloquium (DMK)*, 2022.

[58] C. C. Rheinländer, F. Lauer, and N. Wehn. A New Method for Predictive Checkpointing in Transiently-Powered IoT Sensor Devices with Thermal Energy Harvesting. In *2023 10th IEEE Conference on Technologies for Sustainability (SusTech 2023)*, pages 269–274, 2023.

[59] F. Lauer, M. Schöffel, C. C. Rheinländer, and N. Wehn. Exploration of Thermoelectric Energy Harvesting for Secure, TLS-Based Industrial IoT Nodes. In B. Tekinerdogan, Y. Wang, and L.-J. Zhang, editors, *Internet of Things – ICIOT 2022*, pages 92–107, Cham, 2023. Springer Nature Switzerland.

[60] B. Ransford, J. Sorber, and K. Fu. Mementos: System Support for Long-running Computation on RFID-scale Devices. *SIGARCH Comput. Archit. News*, 39(1):159–170, Mar. 2011.

[61] H. Jayakumar, A. Raha, and V. Raghunathan. QUICKRECALL: A Low Overhead HW/SW Approach for Enabling Computations across Power Cycles in Transiently Powered Computers. In *2014 27th International Conference on VLSI Design and 2014 13th International Conference on Embedded Systems*, pages 330–335, Jan 2014.

[62] D. Balsamo, A. S. Weddell, G. V. Merrett, B. M. Al-Hashimi, D. Brunelli, and L. Benini. Hibernus: Sustaining Computation During Intermittent Supply for Energy-Harvesting Systems. *IEEE Embedded Systems Letters*, 7(1):15–18, March 2015.

[63] M. Magno, L. Spadaro, J. Singh, and L. Benini. Kinetic energy harvesting: Toward autonomous wearable sensing for Internet of Things. In *2016 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM)*, pages 248–254, June 2016.

[64] D. G. Liebermann, L. Katz, M. D. Hughes, R. M. Bartlett, J. McClements, and I. M. Franks. Advances in the application of information technology to sport performance. *Journal of Sports Sciences*, 20(10):755–769, 2002. PMID: 12363293.

[65] M. Gasser. Motion Estimation Using Inertial Sensor Technology with Applications to Sporting Exercises. Technical report, European Athletics Innovation Awards 2014, 2014.

[66] S. Derlien, B. Böhme, L. Leistritz, and U. Smolenski. Validitätsuntersuchung zum neuen, innovativen Ganganalysesystem RehaWatch von Hasomed. *Manuelle Medizin*, 48:254–259, 08 2010.

[67] M. Iosa, P. Picerno, S. Paolucci, and G. Morone. Wearable inertial sensors for human movement analysis. *Expert Review of Medical Devices*, 13(7):641–659, 2016. PMID: 27309490.

[68] S. Šprager and M. Juric. Inertial Sensor-Based Gait Recognition: A Review. *Sensors*, 15:22089–22127, 09 2015.

[69] U. Jensen, P. Kugler, M. Ring, and B. M. Eskofier. Approaching the accuracy–cost conflict in embedded classification system design. *Pattern Analysis and Applications*, 19(3):839–855, 2016.

[70] U. Ghoshdastider, R. Viga, and M. Kraft. Experimental evaluation of a pairwise broadcast synchronization in a low-power Cyber-physical system. In *Wireless Sensors and Sensor Networks (WiSNet), 2015 IEEE Topical Conference on*, pages 50–52, Jan 2015.

[71] A. Bideaux, B. Zimmermann, S. Hey, and W. Stork. Synchronization in wireless biomedical-sensor networks with Bluetooth Low Energy. *Current Directions in Biomedical Engineering*, page 73–76, September 2015.

[72] S. Sridhar, P. Misra, and J. Warrior. CheepSync: A Time Synchronization Service for Resource Constrained Bluetooth Low Energy Advertisers. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, IPSN '15, pages 364–365, New York, NY, USA, 2015. ACM.

[73] S. Sridhar, P. Misra, G. Gill, and J. Warrior. Cheepsync: a time synchronization service for resource constrained bluetooth le advertisers. *Communications Magazine, IEEE*, 54(1):136–143, January 2016.

[74] M. Frustaci, P. Pace, and G. Aloi. Securing the IoT world: Issues and perspectives. In *2017 IEEE Conference on Standards for Communications and Networking (CSCN)*, pages 246–251, Sep. 2017.

[75] S. P. Gochhayat, C. Lal, L. Sharma, D. P. Sharma, D. Gupta, J. A. M. Saucedo, and U. Kose. Reliable and secure data transfer in IoT networks. *Wireless Networks*, Jun 2019.

[76] M. Frustaci, P. Pace, G. Aloi, and G. Fortino. Evaluating Critical Security Issues of the IoT World: Present and Future Challenges. *IEEE Internet of Things Journal*, 5(4):2483–2495, Aug 2018.

[77] J. King and A. I. Awad. A distributed security mechanism for resource-constrained IoT devices. *Informatica*, 40(1), 2016.

[78] U. Banerjee, C. Juvekar, S. H. Fuller, and A. P. Chandrakasan. eeDTLS: Energy-Efficient Datagram Transport Layer Security for the Internet of Things. In *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*, pages 1–6, Dec 2017.

[79] A. Haroon, S. Akram, M. A. Shah, and A. Wahid. E-Lithe: A Lightweight Secure DTLS for IoT. In *2017 IEEE 86th Vehicular Technology Conference (VTC-Fall)*, pages 1–5, Sep. 2017.

[80] S. Sahoo, S. S. Sahoo, P. Maiti, B. Sahoo, and A. K. Turuk. A Lightweight Authentication Scheme for Cloud-Centric IoT Applications. In *2019 6th International Conference on Signal Processing and Integrated Networks (SPIN)*, pages 1024–1029, March 2019.

[81] P. Kaňuch and D. Macko. Optimizing Energy Efficiency of Secured IoT Communication by OpenHip. In *2019 42nd International Conference on Telecommunications and Signal Processing (TSP)*, pages 174–177, July 2019.

[82] O. Gulbudak and E. Santi. FPGA-Based Model Predictive Controller for Direct Matrix Converter. *IEEE Transactions on Industrial Electronics*, 63(7):4560–4570, July 2016.

[83] H. Peyrl, J. Liu, and T. Geyer. An FPGA implementation of the fast gradient method for solving the Model Predictive Pulse Pattern Control problem. In *2013 IEEE International Symposium on Sensorless Control for Electrical Drives and Predictive Control of Electrical Drives and Power Electronics (SLED/PRECEDE)*, pages 1–6, Oct 2013.

[84] H. Ayala, R. Sampaio, D. M. Muñoz, C. Llanos, L. Coelho, and R. Jacobi. Nonlinear model predictive control hardware implementation with custom-precision floating point operations. In *2016 24th Mediterranean Conference on Control and Automation (MED)*, pages 135–140, June 2016.

[85] L. D. Tucci, K. O'Brien, M. Blott, and M. D. Santambrogio. Architectural optimizations for high performance and energy efficient Smith-Waterman implementation on FPGAs using OpenCL. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2017*, pages 716–721, March 2017.

[86] F. Kennel, D. Görges, and S. Liu. Energy Management for Smart Grids With Electric Vehicles Based on Hierarchical MPC. *IEEE Transactions on Industrial Informatics*, 9(3):1528–1537, Aug 2013.

[87] V. Kieling. Exploring Transient Computing and Energy Harvesting Techniques Towards Self-Powered Embedded Devices. Master's thesis, University of Kaiserslautern, 2017.

[88] M. Schmidt. *Entwicklung, Validierung und Anwendung eines neuartigen Messsystems auf der Basis von Inertialsensoren zur Ableitung biomechanischer Merkmale im leichtathletischen Sprint*. PhD thesis, Technische Universität Dortmund, 2018.

[89] F. Berkel. *Contributions to Event-triggered and Distributed Model Predictive Control*. PhD thesis, University of Kaiserslautern, 2019.

[90] 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, 2016. SI: {IGCC} 2014.

[91] L. Spadaro, M. Magno, and L. Benini. Poster Abstract: KinetiSee - A Perpetual Wearable Camera Acquisition System with a Kinetic Harvester. In *2016 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–2, April 2016.

[92] M. Magno, G. A. Salvatore, S. Mutter, W. Farrukh, G. Troester, and L. Benini. Autonomous smartwatch with flexible sensors for accurate and continuous mapping of skin temperature. In *2016 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 337–340, May 2016.

[93] J. Wilhelm, S. Blackshire, and M. Lanzone. Energy-Neutral Data Collection Rate Control for IoT Animal Behavior Monitors. *Applied Sciences*, 7(11), 2017.

[94] 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.

[95] G. V. Merrett. Invited - Energy Harvesting and Transient Computing: A Paradigm Shift for Embedded Systems? In *Proceedings of the 53rd Annual Design Automation Conference*, DAC '16, pages 33:1–33:2, New York, NY, USA, 2016. ACM.

[96] M. Magno, D. Kneubühler, P. Mayer, and L. Benini. Micro Kinetic Energy Harvesting for Autonomous Wearable Devices. In *2018 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM)*, pages 105–110, June 2018.

[97] A. Costilla Reyes, A. Abuellil, J. J. Estrada-López, S. Carreon-Bautista, and E. Sánchez-Sinencio. Reconfigurable System for Electromagnetic Energy Harvesting With Inherent Activity Sensing Capabilities for Wearable Technology. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 66(8):1302–1306, Aug 2019.

[98] R. Bolt, M. Magno, T. Burger, A. Romani, and L. Benini. Kinetic AC/DC Converter for Electromagnetic Energy Harvesting in Autonomous Wearable Devices. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 64(12):1422–1426, Dec 2017.

[99] M. Magno and D. Boyle. Wearable Energy Harvesting: From body to battery. In *2017 12th International Conference on Design Technology of Integrated Systems In Nanoscale Era (DTIS)*, pages 1–6, April 2017.

[100] A. Zhang and L. Zhu. A Promising Way of Energy Harvesting for Implantable Medical Devices – Thermoelectric Generator (TEG). In *2021 International Conference on Smart City and Green Energy (ICSCGE)*, pages 22–25, 2021.

[101] 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, Jan 2017.

[102] J. Vondrak, M. Schmidt, A. Proto, M. Penhaker, J. Jargus, and L. Peter. Using Miniature Thermoelectric Generators for Wearable Energy Harvesting. In *2019 4th International Conference on Smart and Sustainable Technologies (SpliTech)*, pages 1–6, 2019.

[103] R. I. Smith and M. L. Johnston. Analysis of Skin-Worn Thermoelectric Generators for Body Heat Energy Harvesting to Power Wearable Devices. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 7158–7161, 2021.

[104] A. Proto, J. Vondrak, M. Schmidt, J. Kubicek, O. M. Gorjani, J. Havlik, and M. Penhaker. A Flexible Thermoelectric Generator Worn on the Leg to Harvest Body Heat Energy and to Recognize Motor Activities: A Preliminary Study. *IEEE Access*, 9:20878–20892, 2021.

[105] L. Sigrist, N. Stricker, D. Bernath, J. Beutel, and L. Thiele. Thermoelectric Energy Harvesting From Gradients in the Earth Surface. *IEEE Transactions on Industrial Electronics*, 67(11):9460–9470, 2020.

[106] N. Nadaf and P. A. Review on Waste Heat Energy Harvesting using TEG: Applications and Enhancements. In *2021 8th International Conference on Smart Computing and Communications (ICSCC)*, pages 334–339, 2021.

[107] A. D. Dos Santos, S. C. de Brito, A. V. Martins, F. F. Silva, and F. Morais. Thermoelectric Energy Harvesting on Rotation Machines for Wireless Sensor Network in Industry 4.0. In *2021 14th IEEE International Conference on Industry Applications (INDUSCON)*, pages 694–697, 2021.

[108] D. Baek, H. G. Lee, and N. Chang. SmartPatch: A Self-Powered and Patchable Cumulative UV Irradiance Meter. *IEEE Design Test*, 36(1):57–64, Feb 2019.

[109] P. Jokic and M. Magno. Powering smart wearable systems with flexible solar energy harvesting. In *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–4, May 2017.

[110] V. Kartsch, S. Benatti, M. Mancini, M. Magno, and L. Benini. Smart Wearable Wristband for EMG based Gesture Recognition Powered by Solar Energy Harvester. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5, May 2018.

[111] A. Gomez, L. Sigrist, M. Magno, L. Benini, and L. Thiele. Dynamic Energy Burst Scaling for Transiently Powered Systems. In *Proceedings of the 2016 Design, Automation & Test in Europe Conference & Exhibition (DATE 2016)*, pages 349–354, Dresden, Mar 2016.

[112] A. P. Sample, D. J. Yeager, P. S. Powledge, A. V. Mamishev, and J. R. Smith. Design of an RFID-Based Battery-Free Programmable Sensing Platform. *IEEE Transactions on Instrumentation and Measurement*, 57(11):2608–2615, 2008.

[113] V. Talla, B. Kellogg, B. Ransford, S. Naderiparizi, S. Gollakota, and J. R. Smith. Powering the next Billion Devices with Wi-Fi. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, CoNEXT '15, New York, NY, USA, 2015. Association for Computing Machinery.

[114] A. Sample and J. R. Smith. Experimental results with two wireless power transfer systems. In *2009 IEEE Radio and Wireless Symposium*, pages 16–18, 2009.

[115] D. Ma, G. Lan, M. Hassan, W. Hu, and S. K. Das. Sensing, Computing, and Communications for Energy Harvesting IoTs: A Survey. *IEEE Communications Surveys Tutorials*, 22(2):1222–1250, 2020.

[116] Y. Kuang, T. Ruan, Z. J. Chew, and M. Zhu. Energy harvesting during human walking to power a wireless sensor node. *Sensors and Actuators A: Physical*, 254:69–77, 2017.

[117] A. R. El-Sayed, K. Tai, M. Biglarbegian, and S. Mahmud. A survey on recent energy harvesting mechanisms. In *2016 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, pages 1–5, May 2016.

[118] A. Ahmed, I. Hassan, M. Hedaya, T. Abo El-Yazid, J. Zu, and Z. L. Wang. Farms of triboelectric nanogenerators for harvesting wind energy: A potential approach towards green energy. *Nano Energy*, 36:21–29, 2017.

[119] E. Zarepour, M. Hassan, C. T. Chou, and A. A. Adesina. Remote Detection of Chemical Reactions Using Nanoscale Terahertz Communication Powered by Pyroelectric Energy Harvesting. In *Proceedings of the Second Annual International Conference on Nanoscale Computing and Communication*, NANOCOM' 15, New York, NY, USA, 2015. Association for Computing Machinery.

[120] M. B. Upama, M. Wright, N. K. Elumalai, M. A. Mahmud, D. Wang, C. Xu, and A. Uddin. High-Efficiency Semitransparent Organic Solar Cells with Non-Fullerene Acceptor for Window Application. *ACS Photonics*, 4(9):2327–2334, 2017.

[121] K. Yoshikawa, H. Kawasaki, W. Yoshida, T. Irie, K. Konishi, K. Nakano, T. Uto, D. Adachi, M. Kanematsu, H. Uzu, and K. Yamamoto. Silicon heterojunction solar cell with interdigitated back contacts for a photoconversion efficiency over 26%. *Nature Energy*, 2(5), 3 2017.

[122] S. Khalid, I. Raouf, A. Khan, N. Kim, and H. S. Kim. A Review of Human-Powered Energy Harvesting for Smart Electronics: Recent Progress and Challenges. *International Journal of Precision Engineering and Manufacturing-Green Technology*, 6(4):821–851, Aug 2019.

[123] A. Rodriguez, D. Balsamo, Z. Luo, S. P. Beeby, G. V. Merrett, and A. S. Weddel. Intermittently-powered energy harvesting step counter for fitness tracking. In *2017 IEEE Sensors Applications Symposium (SAS)*, pages 1–6, March 2017.

[124] A. Rodriguez Arreola, D. Balsamo, A. K. Das, A. S. Weddell, D. Brunelli, B. M. Al-Hashimi, and G. V. Merrett. Approaches to Transient Computing for Energy Harvesting Systems:

A Quantitative Evaluation. In *Proceedings of the 3rd International Workshop on Energy Harvesting & Energy Neutral Sensing Systems*, ENSsys '15, pages 3–8, New York, NY, USA, 2015. ACM.

[125] M. Kim, J. Lee, Y. Kim, and Y. H. Song. An analysis of energy consumption under various memory mappings for FRAM-based IoT devices. In *2018 IEEE 4th World Forum on Internet of Things (WF-IoT)*, pages 574–579, Feb 2018.

[126] H. Jayakumar, A. Raha, J. R. Stevens, and V. Raghunathan. Energy-Aware Memory Mapping for Hybrid FRAM-SRAM MCUs in Intermittently-Powered IoT Devices. *ACM Trans. Embed. Comput. Syst.*, 16(3):65:1–65:23, Apr. 2017.

[127] K. Maeng and B. Lucia. Supporting Peripherals in Intermittent Systems with Just-in-time Checkpoints. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2019, pages 1101–1116, New York, NY, USA, 2019. ACM.

[128] F. Ahmed, C. Kervadec, Y. Le Moullec, G. Tamberg, and P. Annus. Autonomous Wireless Sensor Networks: Implementation of Transient Computing and Energy Prediction for Improved Node Performance and Link Quality. *The Computer Journal*, 62(6):820–837, 10 2018.

[129] D. Balsamo, A. Elboreini, B. M. Al-Hashimi, and G. V. Merrett. Exploring ARM mbed support for transient computing in energy harvesting IoT systems. In *2017 7th IEEE International Workshop on Advances in Sensors and Interfaces (IWASI)*, pages 115–120, 2017.

[130] D. Balsamo, M. Magno, K. Kubara, B. Lazarescu, and G. V. Merrett. Energy Harvesting Meets IoT: Fuelling Adoption of Transient Computing in Embedded Systems. In *2019 IEEE 5th World Forum on Internet of Things (WF-IoT)*, pages 413–417, 2019.

[131] 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 '17, pages 209–219, New York, NY, USA, 2017. ACM.

[132] K. Ma, Y. Zheng, S. Li, K. Swaminathan, X. Li, Y. Liu, J. Sampson, Y. Xie, and V. Narayanan. Architecture exploration for ambient energy harvesting nonvolatile processors. In *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*, pages 526–537, Feb 2015.

[133] F. Su, K. Ma, X. Li, T. Wu, Y. Liu, and V. Narayanan. Nonvolatile processors: Why is it trending? In *Design, Automation Test in Europe Conference Exhibition (DATE), 2017*, pages 966–971, March 2017.

[134] K. Ma, X. Li, H. Liu, X. Sheng, Y. Wang, K. Swaminathan, Y. Liu, Y. Xie, J. Sampson, and V. Narayanan. Dynamic Power and Energy Management for Energy Harvesting Nonvolatile Processor Systems. *ACM Trans. Embed. Comput. Syst.*, 16(4):107:1–107:23, May 2017.

[135] Revolutionizing context save and restore with MSP FRAM microcontrollers. `https://www.ti.com/lit/wp/slay039/slay039.pdf (2022-05)`.

[136] M. Zwerg, A. Baumann, R. Kuhn, M. Arnold, R. Nerlich, M. Herzog, R. Ledwa, C. Sichert, V. Rzehak, P. Thanigai, and B. O. Eversmann. An 82uA/MHz microcontroller with embedded FeRAM for energy-harvesting applications. In *2011 IEEE International Solid-State Circuits Conference*, pages 334–336, 2011.

[137] ARM mbed. `https://os.mbed.com/ (12.05.2022)`.

[138] Z. Ghodsi, S. Garg, and R. Karri. Optimal checkpointing for secure intermittently-powered IoT devices. In *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 376–383, 2017.

[139] H. Jayakumar, A. Raha, and V. Raghunathan. Energy-Aware Memory Mapping for Hybrid FRAM-SRAM MCUs in IoT Edge Devices. In *2016 29th International Conference on VLSI Design and 2016 15th International Conference on Embedded Systems (VLSID)*, pages 264–269, Jan 2016.

[140] K. Maeng, A. Colin, and B. Lucia. Alpaca: Intermittent Execution Without Checkpoints. *Proc. ACM Program. Lang.*, 1(OOPSLA):96:1–96:30, Oct. 2017.

[141] G. Berthou, T. Delizy, K. Marquet, T. Risset, and G. Salagnac. Peripheral state persistence for transiently-powered systems. In *2017 Global Internet of Things Summit (GIoTS)*, pages 1–6, June 2017.

[142] G. Berthou, T. Delizy, K. Marquet, T. Risset, and G. Salagnac. Sytare: A Lightweight Kernel for NVRAM-Based Transiently-Powered Systems. *IEEE Transactions on Computers*, 68(9):1390–1403, Sep. 2019.

[143] P. Singla, S. S. Singh, and S. R. Sarangi. FlexiCheck: An Adaptive Checkpointing Architecture for Energy Harvesting Devices. In *2019 Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 546–551, March 2019.

[144] A. Colin and B. Lucia. Termination Checking and Task Decomposition for Task-based Intermittent Programs. In *Proceedings of the 27th International Conference on Compiler Construction*, CC 2018, pages 116–127, New York, NY, USA, 2018. ACM.

[145] A. Colin, E. Ruppel, and B. Lucia. A Reconfigurable Energy Storage Architecture for Energy-harvesting Devices. *SIGPLAN Not.*, 53(2):767–781, Mar. 2018.

[146] Y. Xu, H. G. Lee, X. Chen, B. Peng, D. Liu, and L. Liang. Puppet: Energy Efficient Task Mapping For Storage-Less and Converter-Less Solar-Powered Non-Volatile Sensor Nodes. In *2018 IEEE 36th International Conference on Computer Design (ICCD)*, pages 226–233, Oct 2018.

[147] B. Lucia and B. Ransford. A Simpler, Safer Programming and Execution Model for Intermittent Systems. *SIGPLAN Not.*, 50(6):575–585, June 2015.

[148] A. Colin and B. Lucia. Chain: Tasks and Channels for Reliable Intermittent Programs. *SIGPLAN Not.*, 51(10):514–530, Oct. 2016.

[149] M. Hicks. Clank: Architectural Support for Intermittent Computation. *SIGARCH Comput. Archit. News*, 45(2):228–240, June 2017.

[150] A. Gomez, L. Sigrist, T. Schalch, L. Benini, and L. Thiele. Wearable, energy-opportunistic vision sensing for walking speed estimation. In *2017 IEEE Sensors Applications Symposium (SAS)*, pages 1–6, March 2017.

[151] J. Van Der Woude and M. Hicks. Intermittent Computation without Hardware Support or Programmer Intervention. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, OSDI'16, page 17–32, USA, 2016. USENIX Association.

[152] A. Hoseinghorban, M. R. Bahrami, A. Ejlali, and M. A. Abam. CHANCE: Capacitor Charging Management Scheme in Energy Harvesting Systems. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 40(3):419–429, 2021.

[153] K. Ganesan, J. San Miguel, and N. Enright Jerger. The What's Next Intermittent Computing Architecture. In *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 211–223, Feb 2019.

[154] S. Z. Khan, R. Kakar, M. M. Alam, Y. L. Moullec, and H. Pervaiz. A Green IoT Node Incorporating Transient Computing, Approximate Computing and Energy/Data Prediction. In *2020 IEEE 17th Annual Consumer Communications & Networking Conference (CCNC)*, pages 1–6, 2020.

[155] MSP430FR5869 Product Information - Texas Instruments. `https://www.ti.com/product/MSP430FR5869 (2023-12)`.

[156] M. Zangeneh and A. Joshi. Design and Optimization of Nonvolatile Multibit 1T1R Resistive RAM. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 22(8):1815–1828, Aug 2014.

[157] D. Balsamo, M. Magno, K. Kubara, B. Lazarescu, and G. V. Merrett. Energy Harvesting Meets IoT: Fuelling Adoption of Transient Computing in Embedded Systems. In *2019 IEEE 5th World Forum on Internet of Things (WF-IoT)*, pages 413–417, April 2019.

[158] C. Wang, N. Chang, Y. Kim, S. Park, Y. Liu, H. G. Lee, R. Luo, and H. Yang. Storageless and converter-less maximum power point tracking of photovoltaic cells for a nonvolatile microprocessor. In *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 379–384, Jan 2014.

[159] D. Balsamo, B. J. Fletcher, A. S. Weddell, G. Karatziolas, B. M. Al-Hashimi, and G. V. Merrett. Momentum: Power-neutral Performance Scaling with Intrinsic MPPT for Energy Harvesting Computing Systems. *ACM Trans. Embed. Comput. Syst.*, 17(6):93:1–93:25, Jan. 2019.

[160] D. Balsamo, A. Das, A. S. Weddell, D. Brunelli, B. M. Al-Hashimi, G. V. Merrett, and L. Benini. Graceful Performance Modulation for Power-Neutral Transient Computing Systems. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 35(5):738–749, May 2016.

[161] Y. Sun, Z. Yuan, Y. Liu, X. Li, Y. Wang, Q. Wei, Y. Wang, V. Narayanan, and H. Yang. Maximum Energy Efficiency Tracking Circuits for Converter-Less Energy Harvesting Sensor Nodes. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 64(6):670–674, June 2017.

[162] J. Hester, T. Jia, and J. Gu. Holistic Energy Management with μProcessor Co-Optimization in Fully Integrated Battery-Less IoTs. In *2018 31st IEEE International System-on-Chip Conference (SOCC)*, pages 7–12, Sep. 2018.

[163] Z. Gong, K. Qiu, D. Zhou, W. Chen, Y. Xu, X. Shi, and Y. Liu. Retention state-enabled and progress-driven energy management for self-powered nonvolatile processors. In *2017 IEEE 23rd International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, pages 1–8, Aug 2017.

[164] Kinetron. Kinetron. `https://www.kinetron.eu/ (13.01.2023)`, 2019.

[165] Matrix Industries. Prometheus — MATRIX: Self-Powered Solutions. `https://www.matrixindustries.com/prometheus (13.01.2023)`, 2023.

[166] Maxim Integrated. High-Precision I2C Temperature Sensor on a Battery Power Budge. `https://www.stg-maximintegrated.com/en/products/sensors/MAX31889.html (13.01.2023)`, 2022.

[167] STMicroelectronics. STM32L4 - ARM Cortex-M4 ultra-low-power MCUs - STMicroelectronics. `https://www.st.com/en/microcontrollers-microprocessors/stm32l4-series.html (13.01.2023)`, 2023.

[168] A. Baca. Computer-science based feedback systems on sports performance. *International Journal of Computer Science in Sport*, 2(1):20–30, 2003.

[169] F. Fuss, A. Subic, and R. Mehta. The impact of technology on sport — new frontiers. *Sports Technology*, 1:1 – 2, 06 2008.

[170] J. B. Lee, R. B. Mellifont, and B. J. Burkett. The use of a single inertial sensor to identify stride, step, and stance durations of running gait. *Journal of Science and Medicine in Sport*, 13(2):270 – 273, 2010.

[171] P. Picerno, V. Camomilla, and L. Capranica. Countermovement jump performance assessment using a wearable 3D inertial measurement unit. *Journal of Sports Sciences*, 29(2):139–146, 2011. PMID: 21120742.

[172] M. Patterson and B. Caulfield. A method for monitoring reactive strength index. *Procedia Engineering*, 2(2):3115 – 3120, 2010. The Engineering of Sport 8 - Engineering Emotion.

[173] E. Munoz Diaz, O. Heirich, M. Khider, and P. Robertson. Optimal sampling frequency and bias error modeling for foot-mounted IMUs. In *International Conference on Indoor Positioning and Indoor Navigation*, pages 1–9, 2013.

[174] S. Taherian, M. Pias, R. Harle, G. Coulouris, S. Hay, J. Cameron, J. Lasenby, G. Kuntze, I. Bezodis, G. Irwin, and D. Kerwin. Profiling sprints using on-body sensors. In *2010 8th IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, pages 444–449, 2010.

[175] MATLAB - Mathworks. `https://www.mathworks.com/ (2021-01)`.

[176] P. Blank, S. Hofmann, M. Kulessa, and B. Eskofier. miPod 2: a new hardware platform for embedded real-time processing in sports and fitness applications. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, UbiComp '16, page 881–884, New York, NY, USA, 09 2016. Association for Computing Machinery.

[177] J. Elson, L. Girod, and D. Estrin. Fine-grained Network Time Synchronization Using Reference Broadcasts. *SIGOPS Oper. Syst. Rev.*, 36(SI):147–163, Dec. 2002.

[178] S. Ganeriwal, R. Kumar, and M. B. Srivastava. Timing-sync Protocol for Sensor Networks. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems*, SenSys '03, pages 138–149, New York, NY, USA, 2003. ACM.

[179] M. Maróti, B. Kusy, G. Simon, and A. Lédeczi. The Flooding Time Synchronization Protocol. In *Proceedings of the 2Nd International Conference on Embedded Networked Sensor Systems*, SenSys '04, pages 39–49, New York, NY, USA, 2004. ACM.

[180] P. Sommer and R. Wattenhofer. Gradient clock synchronization in wireless sensor networks. In *Information Processing in Sensor Networks, 2009. IPSN 2009. International Conference on*, pages 37–48, April 2009.

[181] K. Noh, E. Serpedin, and K. Qaraqe. A New Approach for Time Synchronization in Wireless Sensor Networks: Pairwise Broadcast Synchronization. *Wireless Communications, IEEE Transactions on*, 7(9):3318–3322, September 2008.

[182] E. Foxlin. Pedestrian tracking with shoe-mounted inertial sensors. *IEEE Computer Graphics and Applications*, 25(6):38–46, 2005.

[183] B. Mariani, C. Hoskovec, S. Rochat, C. Büla, J. Penders, and K. Aminian. 3D gait assessment in young and elderly subjects using foot-worn inertial sensors. *Journal of Biomechanics*, 43(15):2999 – 3006, 2010.

[184] A. Baca and P. Kornfeind. Rapid Feedback Systems for Elite Sports Training. *IEEE Pervasive Computing*, 5(4):70–76, 2006.

[185] W. Sands. Measurement issues with elite athletes. *Sports Technology*, 1:101 – 104, 07 2008.

[186] Technische Universität Dortmund - Institut für Sport und Sportwissenschaft. `https://www.sport.tu-dortmund.de/sport/index.php (2021-01)`.

[187] Technische Universität Dortmund - SpoSeNs project. `https://www.sport.tu-dortmund.de/sport/index.php?page=show_projekt&id=52 (2021-01)`.

[188] C. C. Rheinländer. Conception and Development of a Flexible Multi Sensor Ultra Low Power Wireless Sensor Network using Bluetooth Low Energy. Master's thesis, Fachhochschule (FH) Kaiserslautern, February 2014.

[189] E. Bergamini, P. Picerno, H. Pillet, F. Natta, P. Thoreux, and V. Camomilla. Estimation of temporal parameters during sprint running using a trunk-mounted inertial measurement unit. *Journal of Biomechanics*, 45(6):1123 – 1126, 2012.

[190] B. Purcell, J. Channells, D. James, and R. Barrett. Use of accelerometers for detecting foot-ground contact time during running. *Proc SPIE*, 01 2006.

[191] M. Coh. Biomechanical Analysis of Colin Jackson's Hurdle Clearance Technique. *New Studies in Athletics*, pages 37–45, 2003.

[192] B. Gajer, C. Hanon, and C. Thépaut-Mathieu. Velocity and stride parameters in the 400 Metres. *New Stud Athl*, 22:39–46, 01 2007.

[193] J. Slawinski, N. Termoz, G. Rabita, G. Guilhem, S. Dorel, J.-B. Morin, and P. Samozino. How 100-m event analyses improve our understanding of world-class men's and women's sprint performance. *Scandinavian Journal of Medicine and Science in Sports*, 10 2015.

[194] U. Ghoshdastider, R. Viga, and M. Kraft. Wireless time synchronization of a collaborative brain-computer-interface using bluetooth low energy. In *SENSORS, 2014 IEEE*, pages 2250–2254, Nov 2014.

[195] M. J. Joachim R. Uffelmann, Peter Wienzek. *IO-Link. Brückentechnologie für Industrie 4.0.* Deutscher Industrieverlag, Essen, 2. auflage edition, 2018.

[196] IO-Link Community. IO-Link Wireless Exposé. `https://io-link.com/share/Downloads/At-a-glance/IO-Link_Wireless_Expos%C3%A9_eng_2018.pdf (26.02.2021)`, 2018.

[197] C. C. Rheinländer. SyRealNET Node - Documentation v1.0. Technical report, University of Kaiserslautern (TU-KL), Division of Microelectronic Systems Design (EMS), July 2020.

[198] 3GPP. 3rd Generation Partnership Project - A Global Initiative. `https://www.3gpp.org/ (09.03.2021)`, 2021.

[199] LoRa Alliance. `https://lora-alliance.org/ (18.03.2021)`, 2021.

[200] Sigfox - The Global Communications Service Provider for the Internet of Things (IoT). `https://www.sigfox.com/en (18.03.2021)`, 2021.

[201] V. Hassija, V. Chamola, V. Saxena, D. Jain, P. Goyal, and B. Sikdar. A Survey on IoT Security: Application Areas, Security Threats, and Solution Architectures. *IEEE Access*, 7:82721–82743, 2019.

[202] D. Diaz-Sanchez, A. Marin-Lopez, F. Almenarez, P. Arias, and R. S. Sherratt. TLS/PKI Challenges and Certificate Pinning Techniques for IoT and M2M Secure Communications. *IEEE Communications Surveys Tutorials*, pages 1–1, 2019.

[203] M. Suárez-Albela, T. M. Fernández-Caramés, P. Fraga-Lamas, and L. Castedo. A Practical Performance Comparison of ECC and RSA for Resource-Constrained IoT Devices. In *2018 Global Internet of Things Summit (GIoTS)*, pages 1–6, June 2018.

[204] A. H. Gerez, K. Kamaraj, R. Nofal, Y. Liu, and B. Dezfouli. Energy and Processing Demand Analysis ofTLS Protocol in Internet of Things Applications. In *2018 IEEE International Workshop on Signal Processing Systems (SiPS)*, pages 312–317, Oct 2018.

[205] I. Ud Din, M. Guizani, S. Hassan, B. Kim, M. Khurram Khan, M. Atiquzzaman, and S. H. Ahmed. The Internet of Things: A Review of Enabled Technologies and Future Challenges. *IEEE Access*, 7:7606–7640, 2019.

[206] I. Ud Din, M. Guizani, B. Kim, S. Hassan, and M. Khurram Khan. Trust Management Techniques for the Internet of Things: A Survey. *IEEE Access*, 7:29763–29787, 2019.

[207] S. Oh and Y. Kim. Security Requirements Analysis for the IoT. In *2017 International Conference on Platform Technology and Service (PlatCon)*, pages 1–6, Feb 2017.

[208] TLS Working Group. Compact TLS 1.3 - Draft. `https://tools.ietf.org/html/draft-rescorla-tls-ctls-02 (27.08.2019)`, 2019.

[209] M. Vučinić, B. Tourancheau, F. Rousseau, A. Duda, L. Damon, and R. Guizzetti. OSCAR: Object security architecture for the Internet of Things. In *Proceeding of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks 2014*, pages 1–10, June 2014.

[210] S. Raza, H. Shafagh, K. Hewage, R. Hummen, and T. Voigt. Lithe: Lightweight Secure CoAP for the Internet of Things. *IEEE Sensors Journal*, 13(10):3711–3720, Oct 2013.

[211] M. Ammar, G. Russello, and B. Crispo. Internet of Things: A survey on the Security of IoT Frameworks. *Journal of Information Security and Applications*, 38:8–27, 2018.

[212] AWS IoT Developer Guide. Security and Identity for AWS IoT. `https://docs.aws.amazon.com/iot/latest/developerguide/iot-security-identity.html (27.08.2019)`, 2019.

[213] Cloud IoT Core. Requirements. `https://cloud.google.com/iot/docs/requirements (27.08.2019)`, 2019.

[214] IPv6 over Low power WPAN (6lowpan). `https://datatracker.ietf.org/wg/6lowpan/documents/ (18.03.2021)`, 2021.

[215] IEEE 802.15.4. `https://www.ieee802.org/15/pub/TG4.html (18.03.2021)`, 2021.

[216] R. Tabish, A. Ben Mnaouer, F. Touati, and A. M. Ghaleb. A comparative analysis of BLE and 6LoWPAN for U-HealthCare applications. In *2013 7th IEEE GCC Conference and Exhibition (GCC)*, pages 286–291, 2013.

[217] M. Siekkinen, M. Hiienkari, J. K. Nurminen, and J. Nieminen. How low energy is bluetooth low energy? Comparative measurements with ZigBee/802.15.4. In *2012 IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*, pages 232–237, April 2012.

[218] E. Morin, M. Maman, R. Guizzetti, and A. Duda. Comparison of the Device Lifetime in Wireless Networks for the Internet of Things. *IEEE Access*, 5:7097–7114, 2017.

[219] P. Trelsmo, P. Di Marco, P. Skillermark, R. Chirikov, and J. Ostman. Evaluating IPv6 Connectivity for IEEE 802.15.4 and Bluetooth Low Energy. In *2017 IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*, pages 1–6, March 2017.

[220] Nordic Semiconductor. `https://www.nordicsemi.com/ (18.03.2021)`, 2021.

[221] RFC 7668 - IPv6 over BLUETOOTH(R) Low Energy. `https://tools.ietf.org/html/rfc7668 (18.03.2021)`, 2021.

[222] H. W. van der Westhuizen and G. P. Hancke. Practical Comparison between COAP and MQTT - Sensor to Server level. In *2018 Wireless Advanced (WiAd)*, pages 1–6, June 2018.

[223] S. N. Swamy, D. Jadhav, and N. Kulkarni. Security threats in the application layer in IoT applications. In *2017 International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)*, pages 477–480, Feb 2017.

[224] Eclipse Mosquitto. `https://mosquitto.org/ (18.03.2021)`, 2021.

[225] Raspberry Pi - Small Single-Board Computer. `https://www.raspberrypi.org/ (18.03.2021)`, 2021.

[226] Nordic Semiconductor - nRF5 SDK v16.0.0 - Router Setup. `https://infocenter.nordicsemi.com/index.jsp?topic=%2Fsdk_nrf5_v16.0.0%2Fiot_ug_border_router.html&cp=7_5_0_5_5_0 (18.03.2021)`, 2021.

[227] Nordic Semiconductor - nRF5 SDK v16.0.0 - MQTT. `https://infocenter.nordicsemi.com/index.jsp?topic=%2Fsdk_nrf5_v16.0.0%2Fiot_sdk_app_mqtt.html&cp=7_5_0_4_2_4_6 (18.03.2021)`, 2021.

[228] Lightweight IP stack. `https://www.nongnu.org/lwip/2_1_x/index.html (11.03.2021)`, 2021.

[229] Mbed TLS - A C library implementing cryptographic primitives, X.509 certificate manipulation and the SSL/TLS and DTLS protocols. `https://www.trustedfirmware.org/projects/mbed-tls/ (11.03.2021)`, 2021.

[230] K. MacKay. micro-ecc: A small ECDH and ECDSA implementation for 32-bit microcontrollers. `http://kmackay.ca/micro-ecc/ (11.03.2021)`, 2021.

[231] E. Nurvitadhi, G. Venkatesh, J. Sim, D. Marr, R. Huang, J. Ong Gee Hock, Y. T. Liew, K. Srivatsan, D. Moss, S. Subhaschandra, and G. Boudoukh. Can FPGAs Beat GPUs in Accelerating Next-Generation Deep Neural Networks? In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, FPGA '17, pages 5–14, New York, NY, USA, 2017. ACM.

[232] G. A. Constantinides. Tutorial paper: Parallel architectures for model predictive control. In *In: Proc. of the European Control Conference 2009*, pages 138–143, 2009.

[233] G. Knagge, A. Wills, A. Mills, and B. Ninness. ASIC and FPGA implementation strategies for Model Predictive Control. In *2009 European Control Conference (ECC)*, pages 144–149, 2009.

[234] J. L. Jerez, G. A. Constantinides, E. C. Kerrigan, and K.-V. Ling. Parallel MPC for Real-Time FPGA-based Implementation. *IFAC Proceedings Volumes*, 44(1):1338 – 1343, 2011. 18th IFAC World Congress.

[235] A. G. Wills, G. Knagge, and B. Ninness. Fast Linear Model Predictive Control Via Custom Integrated Circuit Architecture. *IEEE Transactions on Control Systems Technology*, 20(1):59–71, Jan 2012.

[236] Xilinx. Introduction to FPGA Design with Vivado HLS UG998 (v1.1). `https://www.xilinx.com/support/documentation/sw_manuals/ug998-vivado-intro-fpga-design-hls.pdf (25.02.2021)`, 22.01.2019.

[237] A. Suardi, E. C. Kerrigan, and G. A. Constantinides. Fast FPGA prototyping toolbox for embedded optimization. In *2015 European Control Conference (ECC)*, pages 2589–2594, July 2015.

[238] B. Käpernick, S. Süß, E. Schubert, and K. Graichen. A synthesis strategy for nonlinear model predictive controller on FPGA. In *2014 UKACC International Conference on Control (CONTROL)*, pages 662–667, July 2014.

[239] S. Lucia, D. Navarro, Ó. Lucía, P. Zometa, and R. Findeisen. Optimized FPGA Implementation of Model Predictive Control for Embedded Systems Using High-Level Synthesis Tool. *IEEE Transactions on Industrial Informatics*, 14(1):137–145, Jan 2018.

[240] S. Longo, E. C. Kerrigan, K. V. Ling, and G. A. Constantinides. Parallel move blocking Model Predictive Control. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 1239–1244, Dec 2011.

[241] W. Xi, L. Dewei, and X. Yugeng. A mixed-integer quadratic programming solver based on GPU. In *2015 34th Chinese Control Conference (CCC)*, pages 2686–2691, July 2015.

[242] Y. Gang and L. Mingguang. Acceleration of MPC using graphic processing unit. In *Proceedings of 2012 2nd International Conference on Computer Science and Network Technology*, pages 1001–1004, Dec 2012.

[243] E. Smith, J. Gondzio, and J. Hall. GPU Acceleration of the Matrix-Free Interior Point Method. In R. Wyrzykowski, J. Dongarra, K. Karczewski, and J. Waśniewski, editors, *Parallel Processing and Applied Mathematics*, pages 681–689, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[244] J. Fowers, K. Ovtcharov, K. Strauss, E. S. Chung, and G. Stitt. A High Memory Bandwidth FPGA Accelerator for Sparse Matrix-Vector Multiplication. In *2014 IEEE 22nd Annual International Symposium on Field-Programmable Custom Computing Machines*, pages 36–43, May 2014.

[245] S. Kestur, J. D. Davis, and O. Williams. BLAS Comparison on FPGA, CPU and GPU. In *2010 IEEE Computer Society Annual Symposium on VLSI*, pages 288–293, July 2010.

[246] L. Yu, A. Goldsmith, and S. Di Cairano. Efficient Convex Optimization on GPUs for Embedded Model Predictive Control. In *Proceedings of the General Purpose GPUs*, GPGPU-10, pages 12–21, New York, NY, USA, 2017. ACM.

[247] S. Richter, T. Geyer, and M. Morari. Resource-Efficient Gradient Methods for Model Predictive Pulse Pattern Control on an FPGA. *IEEE Transactions on Control Systems Technology*, 25(3):828–841, May 2017.

[248] J. Vallone, T. Geyer, and E. R. Rohr. FPGA implementation of model predictive direct current control. In *2016 IEEE Energy Conversion Congress and Exposition (ECCE)*, pages 1–8, Sept 2016.

[249] J. L. Jerez, G. A. Constantinides, and E. C. Kerrigan. Towards a fixed point QP solver for predictive control. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 675–680, Dec 2012.

[250] J. L. Jerez, P. J. Goulart, S. Richter, G. A. Constantinides, E. C. Kerrigan, and M. Morari. Embedded Predictive Control on an FPGA using the Fast Gradient Method. In *2013 European Control Conference (ECC)*, pages 3614–3620, July 2013.

[251] A. Mills, A. G. Wills, S. R. Weller, and B. Ninness. Implementation of linear model predictive control using a field-programmable gate array. *IET Control Theory Applications*, 6(8):1042–1054, May 2012.

[252] L. Vesely and Z. Mynar. Model predictive control of SPMSM based on FPGA and processor. In *2017 IEEE International Conference on Industrial Technology (ICIT)*, pages 324–329, March 2017.

[253] I. McInerney, G. A. Constantinides, and E. C. Kerrigan. A Survey of the Implementation of Linear Model Predictive Control on FPGAs. *IFAC-PapersOnLine*, 51(20):381 – 387, 2018. 6th IFAC Conference on Nonlinear Model Predictive Control NMPC 2018.

[254] J. L. Jerez, G. A. Constantinides, and E. C. Kerrigan. An FPGA Implementation of a Sparse Quadratic Programming Solver for Constrained Predictive Control. In *Proceedings of the 19th ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, FPGA '11, pages 209–218, New York, NY, USA, 2011. ACM.

[255] B. Khusainov, E. Kerrigan, A. Suardi, and G. Constantinides. Nonlinear predictive control on a heterogeneous computing platform. *Control Engineering Practice*, 78:105 – 115, 2018.

[256] D. Boland and G. A. Constantinides. Optimising Memory Bandwidth Use for Matrix-Vector Multiplication in Iterative Methods. In P. Sirisuk, F. Morgan, T. El-Ghazawi, and H. Amano, editors, *Reconfigurable Computing: Architectures, Tools and Applications*, pages 169–181, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

[257] J. L. Jerez, K. . Ling, G. A. Constantinides, and E. C. Kerrigan. Model predictive control for deeply pipelined field-programmable gate array implementation: algorithms and circuitry. *IET Control Theory Applications*, 6(8):1029–1041, May 2012.

[258] F. Berkel, S. Caba, J. Bleich, and S. Liu. A modeling and distributed MPC approach for water distribution networks. *Control Engineering Practice*, 81:199 – 206, 2018.

[259] A. Wills, A. Mills, and B. Ninness. FPGA Implementation of an Interior-Point Solution for Linear Model Predictive Control*. *IFAC Proceedings Volumes*, 44(1):14527–14532, 2011. 18th IFAC World Congress.

[260] E. N. Hartley, J. L. Jerez, A. Suardi, J. M. Maciejowski, E. C. Kerrigan, and G. A. Constantinides. Predictive Control Using an FPGA With Application to Aircraft Control. *IEEE Transactions on Control Systems Technology*, 22(3):1006–1017, 2014.

[261] H. A. Shukla, B. Khusainov, E. C. Kerrigan, and C. N. Jones. Software and Hardware Code Generation for Predictive Control Using Splitting Methods. *IFAC-PapersOnLine*, 50(1):14386 – 14391, 2017. 20th IFAC World Congress.

[262] Modbus. The Modbus Organization. `https://modbus.org/ (25.02.2021)`, 2021.

[263] Digilent. ZedBoard Zynq-7000 ARM/FPGA SoC Development Board - Digilent. `https://store.digilentinc.com/zedboard-zynq-7000-arm-fpga-soc-development-board/ (25.02.2021)`, 2021.

[264] R. Landauer. Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[265] MIT Energy Initiative. Energy-efficient computing - Enabling smaller, lighter, faster computers. `https://energy.mit.edu/news/energy-efficient-computing/ (10.05.2023)`, June 2013.

[266] Statista. Global E-Waste - Statistics & Facts. `https://www.statista.com/topics/3409/electronic-waste-worldwide/#topicOverview (26.04.2023)`, February 2023.

[267] Majeed Ahmad. Wireless IoT Designs and the Evolution of RF Modules. https://www.5gtechnologyworld.com/wireless-iot-designs-and-the-evolution-of-rf-modules/ (31.03.2022), May 2019.

[268] RUSS ARENSMAN. Mixed-signal designers find the right mix. https://www.edn.com/mixed-signal-designers-find-the-right-mix/ (31.03.2022), June 2003.

[269] Bluetooth SIG. Core Specification 4.0 – Bluetooth Technology Website. https://www.bluetooth.com/specifications/specs/core-specification-4-0/ (31.03.2022), June 2010.

[270] R. Heydon. *Bluetooth Low Energy: The Developer's Handbook*. Pearson, 2012.

[271] Nordic Semiconductor. nRF8001, Single-chip Bluetooth low energy solution, Product Specification 1.0 [referring to v0.9.2]. https://infocenter.nordicsemi.com/pdf/nRF8001_PS_v1.0.pdf (31.03.2022), May 2011.

[272] Peter Clarke. Nordic launches Bluetooth LE slave chip. https://www.eetimes.com/nordic-launches-bluetooth-le-slave-chip/ (31.03.2022), 2011.

[273] S. Nagod1 and D. S. Halse. Evolution of MEMS Technology. *International Research Journal of Engineering and Technology (IRJET)*, 4:137–140, 2017.

[274] D. Eddy and D. Sparks. Application of MEMS technology in automotive sensors and actuators. *Proceedings of the IEEE*, 86(8):1747–1755, 1998.

[275] G. Langfelder, M. Bestetti, and M. Gadola. Silicon MEMS inertial sensors evolution over a quarter century. *Journal of Micromechanics and Microengineering*, 31(8):084002, jul 2021.

[276] i-micronews. The MEMS success story: a decade of evolving strategies leading to a bright MEMS future. https://www.i-micronews.com/the-mems-success-story-a-decade-of-evolving-strategies-leading-to-a-bright-mems-future-an-interview-with-bosch-sensortec-stmicroelectronics-and-yole-developpement/ (31.03.2022), October 2021.

[277] Example of a MEMS layout: A torsional ratcheting actuator. https://sites.google.com/site/anishajaindesign/ (18.05.2021).

[278] P. D. Mitcheson. Energy harvesting for human wearable and implantable bio-sensors. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pages 3432–3436, Aug 2010.

[279] QLC Flash - Micron. https://www.micron.com/products/nand-flash/qlc-nand (2020-12).

[280] 3D Flash - Micron. https://www.micron.com/products/nand-flash/3d-nand (2020-12).

[281] N. Gonzales, J. Dinh, D. Lewis, N. Gilbert, B. Pedersen, D. Kamalanathan, J. R. Jameson, and S. Hollmer. An Ultra Low-Power Non-Volatile Memory Design Enabled by Subquantum Conductive-Bridge RAM. In *2016 IEEE 8th International Memory Workshop (IMW)*, pages 1–4, 2016.

[282] TLC Flash - Micron. https://www.micron.com/products/nand-flash/tlc-nand (2020-12).

[283] FRAM - Texas Instruments. http://www.ti.com/lit/pdf/slat151 (2020-12).

[284] FRAM - Cypress. https://www.cypress.com/products/f-ram-nonvolatile-ferroelectric-ram (2020-12).

[285] FRAM - Lapis Tech. https://www.lapis-tech.com/en/semicon/memory/feram.html#TabPage2 (2020-12).

[286] FRAM - Fujitsu. https://www.fujitsu.com/global/documents/products/devices/semiconductor/fram/overview/brochure/AD05-00033-11E.pdf (2020-12).

[287] ReRAM MCU - Panasonic. https://industrial.panasonic.com/ww/products/semiconductors/microcomputers/mn101l/models (2020-12).

[288] CBRAM - Adesto and Dialog Semiconductor. https://www.dialog-semiconductor.com/products/memory/cbram-technology (2020-12).

[289] MRAM - Renesas. https://www.renesas.com/eu/en/products/memory-logic/mrams (2020-12).

[290] MRAM - Everspin. https://www.everspin.com/toggle-mram-technology (2020-12).

[291] J. S. Meena, S. M. Sze, U. Chand, and T.-Y. Tseng. Overview of emerging nonvolatile memory technologies. *Nanoscale Research Letters*, 9(526), 09 2014.

[292] STT-MRAM - Everspin. https://www.everspin.com/spin-transfer-torque-mram-technology (2020-12).

[293] FRAM - Cypress. https://www.cypress.com/file/46186/download (2020-12).

[294] J. R. Jameson, P. Blanchard, C. Cheng, J. Dinh, A. Gallo, V. Gopalakrishnan, C. Gopalan, B. Guichet, S. Hsu, D. Kamalanathan, D. Kim, F. Koushan, M. Kwan, K. Law, D. Lewis, Y. Ma, V. McCaffrey, S. Park, S. Puthenthermadam, E. Runnion, J. Sanchez, J. Shields, K. Tsai, A. Tysdal, D. Wang, R. Williams, M. N. Kozicki, J. Wang, V. Gopinath, S. Hollmer, and M. Van Buskirk. Conductive-bridge memory (CBRAM) with excellent high-temperature retention. In *2013 IEEE International Electron Devices Meeting*, pages 30.1.1–30.1.4, 2013.

[295] J. R. Jameson, J. Dinh, N. Gonzales, S. Hollmer, S. Hsu, D. Kim, F. Koushan, D. Lewis, E. Runnion, J. Shields, A. Tysdal, D. Wang, and V. Gopinath. Towards Automotive Grade Embedded RRAM. In *2018 48th European Solid-State Device Research Conference (ESSDERC)*, pages 58–61, 2018.

[296] D. Görges. Lecture: Model Predictive Control, January 2015.

[297] M. Diehl. *Optimization Algorithms for Model Predictive Control*, pages 1–11. Springer London, London, 2013.

[298] P. Giselsson, M. D. Doan, T. Keviczky, B. De Schutter, and A. Rantzer. Accelerated gradient methods and dual decomposition in distributed model predictive control. *Automatica*, 49(3):829–833, 2013.

[299] R. Scattolini. Architectures for distributed and hierarchical Model Predictive Control - A review. *Journal of Process Control*, 19(5):723 – 731, 2009.

List of Figures

List of Tables

List of Algorithms

Lebenslauf

Persönliche Daten

Name: Carl Christian Rheinländer

Ausbildung

Hochschulstudium:

Bachelor Informationstechnik
Fachhochschule Kaiserslautern
10/2008–08/2012
Bachelor im August 2012

Auslandssemester Bachelor Elektrotechnik
San Diego State University, Kalifornien, USA
08/2011–02/2012

Master Elektrotechnik und Informationstechnik
Fachhochschule Kaiserslautern
10/2012–02/2014
Master im Februar 2014

Beruflicher Werdegang

03/2015–03/2021
Wissenschaftlicher Mitarbeiter
Arbeitsgruppe Entwurf Mikroelektronischer Systeme
Fachbereich Elektrotechnik und Informationstechnik
Technische Universität Kaiserslautern

seit 04/2021
Angestellter als Elektronikentwickler
Hauptabteilung 5, Abteilung 5.5 - Zentrale Elektronik
Technische Universität Kaiserslautern
Seit 09/2022 in Teilzeit (50%)

seit 09/2022
Angestellter als Hard- und Softwareentwickler
in Teilzeit (50%)
Firma Wille Engineering, Hattersheim am Main