

Digital Circuits and Systems for Ultra-Low Power Internet of Things Applications

by

Jacob S. Breiholz

A dissertation presented to

The Faculty of the School of Engineering and Applied Science

University of Virginia

In partial fulfillment of the requirements for the degree

Doctor of Philosophy in Electrical Engineering

December 2020

APPROVAL SHEET

This
Dissertation
is submitted in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

Author: Jacob S. Breiholz

This Dissertation has been read and approved by the examining committee:

Advisor: Benton H. Calhoun

Advisor:

Committee Member: Mircea R. Stan

Committee Member: Bradford Campbell

Committee Member: Joanne Bechta Dugan

Committee Member: Avik Ghosh

Committee Member:

Committee Member:

Accepted for the School of Engineering and Applied Science:



Craig H. Benson, School of Engineering and Applied Science

May 2021

Abstract

The Internet of Things (IoT) is growing at an exponential rate, and some estimate that it will reach 1 trillion devices by 2035. Although IoT devices are functionally diverse, many of them are designed to sense, process, and transmit information. IoT sensing applications are numerous and span many sectors of the global economy including healthcare, manufacturing, agriculture, and infrastructure. The IoT has the potential to produce tremendous amounts of information that can be used to improve human life, but the proliferation of IoT devices is currently being hindered by the challenge of large-scale battery recharge and replacement, which increases the cost of device deployment. In recent years, significant research has focused on lowering the power consumption of IoT sensing devices so that they can operate from ambient harvested energy instead of batteries. This eliminates the need for battery recharge and replacement, and significantly reduces the maintenance costs of large-scale sensor networks. Lower power consumption improves self-powered device lifetime and reliability, and enables operation with smaller energy harvesters in more energy-scarce environments. Ultimately, if the cost of deployment can be reduced below the value of the information being collected, widespread adoption of IoT sensing devices will occur.

This dissertation presents contributions in integrated circuit and self-powered system design for reducing the power consumption of IoT sensing devices and accelerating their adoption. This is accomplished using digital circuits, digital circuit design techniques, and system-level strategies that reduce the power of wireless Transmitter (TX) and digital processing circuits in Ultra-Low Power (ULP) System on Chips (SoCs). Wireless TX and digital processing circuits are both essential components in ULP SoCs, and often dominate the overall system power in many applications. The design of several ULP SoCs, their applications, and integration into larger sensing systems is also presented to demonstrate the technology readiness of self-powered

devices and encourage commercial adoption.

This dissertation presents several research contributions. Chapter 2 presents a lossless data compression accelerator designed to reduce the amount of sensor data that must be transmitted, and thus wireless TX power. Test chip measurements show that this accelerator adds only 4.4 nW processing power overhead, and reduces overall SoC power by 2.9x for an application that samples Electrocardiogram (ECG) data at 360 Hz. Chapter 3 presents a RISC-V microprocessor core implemented using a new performance-scalable version of Dynamic Leakage Suppression (DLS) logic, which enables a continuous power-performance trade-off at power levels below the leakage floor of conventional static CMOS circuits. DLS logic is an emerging method for implementing digital circuits that is significantly lower power than conventional static CMOS, but much slower. This microprocessor core can scale its performance from 6 nW at 11 Hz to 140 nW at 8.2 kHz at a constant 0.6 V supply voltage, and achieves a minimum power consumption of 840 pW across the full supply voltage range. Chapter 4 presents a gate-level multi-threshold technique designed to bridge the power-performance gap between different variants of DLS logic. An array of 8-bit Multiply-Accumulate (MAC) blocks demonstrate that this technique enables up to 6.5x power reduction, and power savings throughout the 10 Hz–1.4 kHz range across the 0.2 V–1.0 V supply voltage range. Chapter 5 presents two ULP SoCs, their IoT sensing systems, and applications. The first SoC implements several health-monitoring applications, and is also integrated into a larger wearable cardiac monitoring system designed for comfortable long-term use. The wearable system is entirely powered by body heat, and consumes just 65 μ W when wirelessly transmitting ECG data to a smartphone. The second SoC is part of a wireless wake-up and control system designed to reduce the electricity usage of miscellaneous electric loads (common household appliances such as power adapters, televisions, and coffee makers) and consumes just 85 nW for this application. Together, these contributions in low-power digital circuit and self-powered system design will help realize the vision of a 1 trillion or more device IoT.

Acknowledgements

First and foremost, I must acknowledge my advisor Benton H. Calhoun. Ben's classes that I attended as an undergraduate were my first introduction to the field of integrated circuit design, and they were by far the most interesting and exciting classes of my undergraduate education. It is likely that I would not have pursued the field further if I did not have such a great first teacher. Ben has been a phenomenal advisor to me during my time in graduate school, and has both inspired and taught me how to be a better version of myself. I am grateful that I had the opportunity to learn from him.

Next I must acknowledge my committee members: Mircea R. Stan, Bradford Campbell, Joanne Bechta Dugan, and Avik Ghosh. I am very grateful for the time and effort that they invested in advising my dissertation, and also that I had the opportunity to learn from many of them as a student in their classes throughout my undergraduate and graduate education. They all contributed greatly to making me the engineer that I am today. I must also acknowledge the funding sources that contributed to this work. These include the NSF NERC ASSIST Center under grant EEC-1160483, and the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Award Number DEEE0008225.

I must also acknowledge the other students who contributed significantly to the work in this dissertation. Farah Yahya and Christopher J. Lukas helped me immensely when I first started graduate school by teaching me the basics of digital circuit design, and also about the research process in general by letting me help out with their projects. They contributed significantly to the design of the sensor data compression accelerator in chapter 2, and also lead the design of the health-monitoring SoC featured in chapter 5. Daniel S. Truesdell contributed significantly to the work presented in chapters 3 and 4 on DLS logic. He also continuously inspired me to be a better researcher throughout graduate school, especially with his incredibly detailed and

visually appealing figures. Without him, much of the work in this dissertation would not have been possible. Luis Lopez Ruiz contributed significantly to the implementation of the wearable cardiac monitoring system presented in chapter 5, and also many earlier versions of this system that are not covered in this dissertation. I am grateful for his help, and that I did not have to deal with all of the challenges of this project by myself.

Next I must acknowledge all of the other incredibly smart and talented students in the Robust Low Power VLSI group that I had the opportunity to work with, and be friends with during graduate school. They made the process much more fun and enjoyable, and I accomplished much more with their help than I would have on my own. In no particular order, their names are: Shuo Li, Henry Bishop, Rishika Agarwala, Sumanth Kamineni, Anjana Dissanayake, Peng Wang, Shourya Gupta, Natalie Ownby, Katy Flynn, Xinjian Liu, Oluseyi Ayorinde, Arijit Banerjee, Ningxi Liu, Harsh Patel, Abhishek Roy, Divya Akella, and He Qi. I must also thank Terry Tigner for all her help over the years with administrative tasks, and her support for all members of the group in general.

Finally, I must thank and acknowledge my family for all their love and support throughout my education. My parents for their emotional, practical, and financial support over the years—they provided me with many opportunities that they did not have in life, and for that I am grateful. I must also acknowledge my wonderful girlfriend Nicole Ulickas. I am so glad that I met her during my time in graduate school, and I am grateful for all her love and support. I look forward to our future together.

Contents

Abstract	iii
Acknowledgements	v
List of Figures	ix
List of Tables	xii
1 Introduction	1
1.1 Motivation	1
1.2 Thesis	5
1.2.1 Reducing Wireless Transmitter Power	6
1.2.2 Reducing Digital Circuit Power	6
1.2.3 Accelerating Proliferation of Self-Powered IoT Sensing Devices	7
1.3 Contributions	7
1.3.1 Ultra-Low Power Sensor Data Compression	7
1.3.2 Ultra-Low Power DVFS RISC-V Microprocessor	8
1.3.3 Multi-Threshold DLS Logic	9
1.3.4 Ultra-Low Power IoT Systems and Applications	9
2 Ultra-Low Power Sensor Data Compression	11
2.1 Introduction	11
2.2 Motivation	13
2.3 Prior Art	14
2.4 A 4.4 nW Lossless Sensor Data Compression Accelerator for 2.9x System Power Reduction in Wireless Body Sensors	18
2.4.1 Contributions	18
2.4.2 Algorithm Definition	18
2.4.3 Algorithm Performance	20
2.4.4 Implementation and Measured Results	22
2.4.5 Conclusion	24
3 Ultra-Low Power DVFS RISC-V Microprocessor	26
3.1 Introduction	26
3.2 Motivation	29
3.3 Prior Art	32
3.4 A 6–140-nW 11-Hz–8.2-kHz DVFS RISC-V Microprocessor Using Scalable Dynamic Leakage Suppression Logic	35
3.4.1 Contributions	35
3.4.2 Scalable Dynamic Leakage Suppression Logic	36

3.4.3	System Architecture	39
3.4.4	Measured Results	39
3.4.5	Conclusion	43
4	Multi-Threshold DLS Logic	45
4.1	Motivation	45
4.2	Prior Art	45
4.3	A Multi-Threshold Technique For Up To 6.5x Power Reduction in Dynamic Leakage Suppression Logic	47
4.3.1	Contributions	47
4.3.2	MT-DLS Logic Implementation	48
4.3.3	Simulated Results	49
4.3.4	Conclusion	51
5	Ultra-Low Power IoT Systems and Applications	53
5.1	Motivation	53
5.2	IoT SoC for Health-Monitoring Applications	54
5.2.1	Free Fall Detection	54
5.2.2	Temperature Sensing	57
5.2.3	Arrhythmia Detection	57
5.2.4	Wearable Cardiac Monitoring System Powered by Body Heat	59
5.3	IoT SoC for MELs Power Management	61
5.3.1	Motivation	61
5.3.2	Contributions	63
5.3.3	System Architecture	64
5.3.4	Measured Results	65
5.4	Conclusion	67
6	Conclusion	68
6.1	Summary of Contributions	68
6.2	Conclusions and Open Problems	71
6.3	Individual Contributions	75
A	Publications	77
A.1	Completed	77
A.2	Planned	78
B	Acronyms	79
C	References	82

List of Figures

1.1	IoT applications are diverse and span many sectors of the global economy. . . .	2
1.2	Block diagram of a typical self-powered SoC targeting IoT sensing applications.	4
2.1	On-chip sensor data compression reduces the amount of data that must be transmitted.	11
2.2	(a) Data framed into 16b words, and (b) variable length data, without the need for framing.	15
2.3	Example binary codes with the prefix property, where no shorter code is the prefix of a longer code.	16
2.4	Block diagram of the LEC compression algorithm. The input is represented as r_i , and the difference between the current input and the predicted input is represented as d_i	16
2.5	Block diagram of the GAS-LEC compression algorithm.	19
2.6	The range and variability of the differences is much less than that of the actual ECG signal. The data is from MIT-BIH record 100 and the DC offset was removed from the ECG waveform to better compare the variability.	21
2.7	The proposed algorithm is effective at compressing 46 ECG records from the MIT-BIH database. Modified lead II (MLII) data was used.	22
2.8	The proposed algorithm is effective at compressing acceleration data for various activities of daily living from the UCI repository. Data from participant 1 was used.	22
2.9	(a) Block diagram of the SoC, showing the compression accelerator is on the TX interface data path. (b) Annotated die photograph of the SoC, showing the compression accelerator (part of control logic).	23
2.10	System-level results, showing the power savings of compression. The SoC power was measured and the TX power was calculated based on the duty cycle and previous measured results for active and sleep power.	25
3.1	Block diagram of a typical clock gate cell.	28
3.2	The normalized I_{DS} as a function of V_{GS} for different V_{TH} devices.	29
3.3	Block diagram of typical cutoff structure implemented using a (a) PMOS header and (b) NMOS footer.	30
3.4	The DLS inverter logic gate.	32
3.5	The DLS inverter in steady state, (a) with the pull-up logic leaking, and (b) with the pull-down logic leaking.	33
3.6	The VTC for a DLS inverter logic gate.	34
3.7	The dual-mode DLS inverter logic gate.	36
3.8	Measured performance with comparison with the state-of-the-art low-power processors.	37
3.9	(a) The SDLS inverter logic gate schematic, and (b) the SDLS inverter logic gate layout compared to a standard static CMOS inverter logic gate.	37

3.10 (a) Transient operation of an SDLS inverter logic gate. (b) The bias voltage effect on steady state voltages and on/off current.	38
3.11 System architecture of the proposed DVFS RISC-V microprocessor.	40
3.12 (a) The standard cells included in the SDLS library, and (b) the measured output of the VSC.	40
3.13 Oscilloscope measurement of runtime DVFS, where the core dithers between two V_{CSEL} modes while computing the Fibonacci sequence. GPIO output bits indicate when a new Fibonacci number is computed and when the total sequence (numbers 1–46) are correct.	41
3.14 The measured power of the core and uncore across DVFS modes at 0.6 V V_{DD}	42
3.15 The measured maximum frequency and energy across DVFS modes at 0.6 V V_{DD}	42
3.16 The power breakdown for the minimum and maximum DVFS modes at 0.6 V V_{DD}	43
3.17 Measured core DVFS tuning range versus V_{DD}	44
3.18 Annotated die photograph of the SDLS RISC-V microprocessor.	44
3.19 Measured performance with comparison with the state-of-the-art low-power processors.	44
4.1 The Forward Body Biased (FBB) DLS inverter logic gate.	46
4.2 Simulated power consumption of an 8-bit MAC block implemented with IO FBB DLS, L_{VT} FBB DLS, and H_{VT} static CMOS at 0.4 V supply voltage.	48
4.3 Block diagram of the 8-bit Multiply-Accumulate (MAC) block.	48
4.4 (a) The DLS standard cell library used in this section. (b) Area comparison between DLS and static CMOS for the 8-bit MAC block.	49
4.5 Simulated power consumption of the MT-DLS 8-bit MAC blocks at 0.4 V supply voltage.	50
4.6 Simulated power consumption and max clock frequency across the 0.2 V–1.0 V supply voltage range for the 8-bit MAC block implemented using L_{VT} FBB DLS logic, IO FBB DLS logic, and MT-DLS logic with 65% L_{VT} FBB logic gates.	51
4.7 Simulated leakage power consumption and maximum clock frequency for the 8-bit MAC blocks implemented with varying percentages of L_{VT} FBB and IO FBB DLS logic gates.	52
4.8 Simulated power reduction for the 8-bit MAC block enabled by MT-DLS across clock frequency and supply voltage. Power reduction is determined by comparison to the MAC block implemented with all L_{VT} FBB logic.	52
5.1 (a) Block diagram of the health-monitoring SoC, and (b) the annotated die photograph.	54
5.2 Flow diagram for the proposed free fall detection application.	56
5.3 (a) Acceleration data captured by the SoC for a free fall. (b) Measured SoC power consumption for the free fall detection application.	57
5.4 (a) ECG data captured by the SoC AFE and ADC. (b) Measured SoC power consumption for the ECG capture application.	58
5.5 The captured ECG signal, showing AFib being detected. The annotations indicate the number of samples between consecutive R peaks.	59

5.6	Block diagram of the proposed three PCB system, and the PCBs integration into the 3D-printed case.	60
5.7	The proposed wearable system is integrated into a custom e-textile shirt, and streams ECG data in real-time using BLE to a smartphone, from which the data is then uploaded to a cloud system.	60
5.8	(a) The system power consumption as a function of the super-capacitor voltage, and (b) the super capacitor voltage as a function of time with no harvesting.	61
5.9	The proposed wireless wake-up and control system for addressing MEL power consumption.	64
5.10	(a) The proposed node-controlling SoC system block diagram, and (b) the SoC die photograph.	65
5.11	State diagram for the example application, and flow chart showing the MEL operating mode transitions	65
5.12	Measured waveform showing wake-up signal, and the GPIO bits that represent the MEL power mode, MEL busy status, and the RF transceiver input	66
5.13	The SoC digital subsystem power consumption as a function of operating frequency and supply voltage.	66
5.14	The combined test setup with the SoC and WuRX, and the measured wakeup signal.	67
6.1	Packet loss between the TX and RX can prevent the data from being decoded correctly.	73

List of Tables

- 1.1 Approximate available power from a 1 cm² energy harvester. 3
- 1.2 Sampling frequency requirements for common IoT sensing applications. 5

- 2.1 Lower Prefix Table 20
- 2.2 Upper Prefix Table 20
- 2.3 Statistics showing the effectiveness of the linear prediction step on ECG and acceleration data. 21
- 2.4 CR comparison for state-of-the-art compression algorithms for ECG and various acceleration data sets. 23
- 2.5 Comparison with state-of-the-art on-chip lossless data compression accelerators. 25

Chapter 1

Introduction

1.1 Motivation

The Internet of Things (IoT) refers to the paradigm of everyday objects becoming connected to the internet through embedded electrical devices. These devices interact with the world through either sensors that collect information about the object or its environment, or actuators that perform some task.¹ They also contain some form of wired or wireless link that connects them to the internet. In a sense, the IoT acts as a bridge between the physical and digital worlds, and thus has the potential to radically change everyday life, just like the internet did. It will do this by providing large amounts of information that will create new insights, leading to better decision making, opportunities for optimization, cost reduction, and improved outcomes.

IoT devices will become prolific when the economic incentive becomes real, and the value of the information collected by the devices is greater than the cost of their deployment.² Large-scale semiconductor companies are predicting that the IoT could reach 1 trillion devices by 2035 due to this macroeconomic principle and expected future cost reduction of IoT system deployment [1]. However, several technological challenges must still be overcome before the IoT can become truly pervasive and scale to over 1 trillion devices. One problem is the feasibility challenge of large-scale battery recharge and replacement. It's not possible for all IoT devices to be connected to the power grid, since that would severely limit the places IoT devices could be deployed. It would also greatly increase the cost of deployment, thus a remote power source

¹This dissertation primarily focuses on IoT devices for sensing information, and generically uses the term IoT device for this.

²Excluding non-economic factors, such as environmental or privacy concerns, which could potentially hinder future IoT growth.

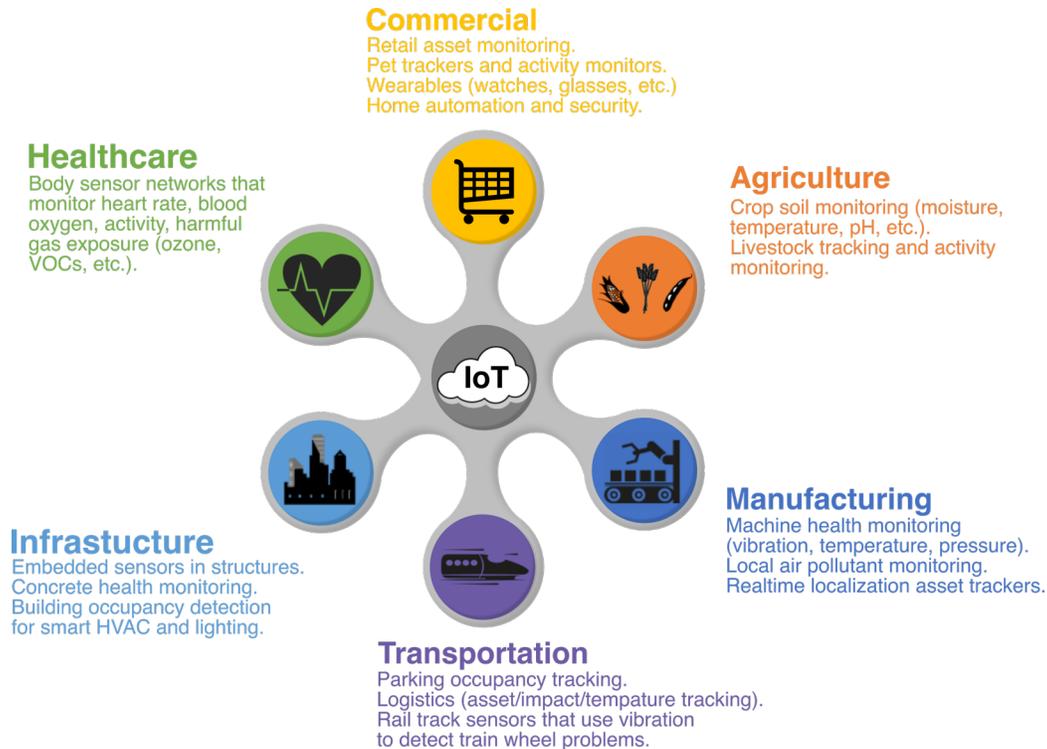


Fig. 1.1. IoT applications are diverse and span many sectors of the global economy.

such as batteries must be used. However, even assuming a conservative 10 year battery lifetime, a 1 trillion device IoT would require approximately 275 million battery replacements or recharges per day to maintain. A less conservative 1 year battery lifetime would require nearly 3 billion battery replacements or recharges per day to maintain.

$$10^{12} \text{ devices} \left(\frac{1 \text{ battery}}{10 \text{ year}} \right) \left(\frac{1 \text{ year}}{365 \text{ days}} \right) \approx 275 \times 10^6 \text{ batteries/day} \quad (1)$$

This represents a large time and financial cost, and has lead researchers to propose making IoT devices self-powered. Self-powered devices rely on harvesting energy from ambient sources such as light, heat, vibration, or Radio Frequency (RF) signals, eliminating the need for battery replacement or recharging. The amount of energy that can be harvested from the environment, however, is often relatively small and can vary greatly depending on the source and ever-changing conditions. Table 1.1 lists the approximate power available from a 1 cm² energy harvester. This creates a strong motivation to lower the power consumption of IoT devices to

the point where they can operate robustly within the limits of energy harvesting.

Table 1.1. Approximate available power from a 1 cm² energy harvester.

Source	Harvested Power		Source
Mechanical			
Human movement	1 – 10	μW	[2]
Machine vibration	10 – 100	μW	[2]
Thermoelectric			
Human body heat	9 – 25	μW	[3]
Machine heat	1 – 10	mW	[2]
Light			
Indoor (bright office)	10	μW	[2]
Outdoor (direct sunlight)	10	mW	[2]
Radio Frequency			
GSM900	36	nW	[4]
GSM1800	84	nW	[4]
3G	12	nW	[4]
WiFi	180	pW	[4]
Other			
Atmospheric moisture	8	μW	[5]
Biochemical (glucose/O ₂)	58	nW	[6]

In addition to being self-powered, it is also highly desirable for IoT devices to have a small form to allow them to be placed anywhere unobtrusively. This imposes strict constraints on both the size of the energy harvester and the size of the energy storage device that can be used. For some applications such as wearables and medical implants, devices can be constrained to just a few cubic millimeters [7, 8]. Common energy harvesters are Photovoltaic (PV) cells for light and Thermoelectric Generators (TEGs) for thermal gradients, and their size is directly proportional to the amount of energy they can harvest. Common energy storage devices are batteries and super capacitors, and their capacity to store energy is also directly proportional to their size. Energy storage devices act as a buffer for fluctuating energy harvesting conditions, so their capacity must be determined according to the variability of the environment and power consumption of the device, such that the device can operate reliably without losing power. These factors further motivate lowering the power consumption of IoT devices, since it permits further reduction in the size of the energy harvester and storage device, which increases the number of places and ways the devices can be deployed.

Another challenge that must be overcome to realize a 1 trillion device IoT is the financial

cost. The cost per device must be low to ensure that the value of the information collected by the device is greater than the cost of the device and its deployment. The lower the cost of the device and its deployment, the more places and ways it will make economic sense to use. IoT devices must therefore be easy to manufacture in large quantities, which motivates integrating as much as possible onto a single silicon substrate, since the infrastructure to cheaply mass produce Integrated Circuits (ICs) already exists. ICs that integrate all of the components needed for a particular system or application are typically referred to as System on Chips (SoCs). Many Ultra-Low Power (ULP) SoCs have been designed for IoT applications [9–14], and they typically contain sensor interfaces for collecting information from the outside world, digital circuits for processing it, memory circuits for storing it, RF circuits for transmitting it, and power management circuits for harvesting energy and regulating it into usable supply voltages for the system. Fig. 1.2 shows a block diagram of a typical self-powered SoC targeting IoT sensing applications. Since SoCs typically integrate all the components needed for a particular system, a complete self-powered IoT sensing device can be created by combining an SoC with an energy harvester and energy storage device. A radio antenna or an off-chip sensor may also be needed depending on the application.

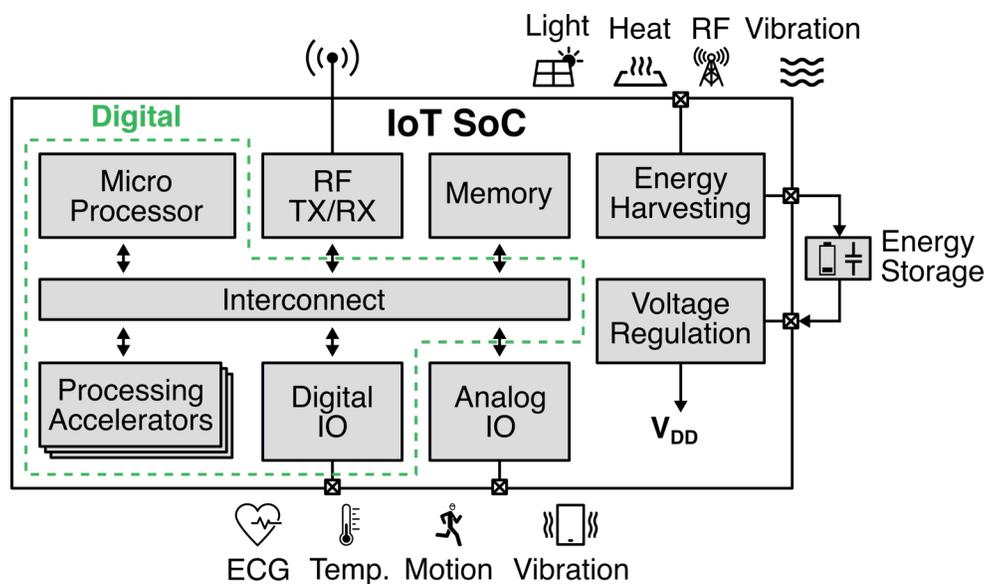


Fig. 1.2. Block diagram of a typical self-powered SoC targeting IoT sensing applications.

Many IoT sensing applications have low sampling frequency requirements in the range of mHz to 100s of Hz, as shown in Table 1.2. This permits aggressive optimization for low power consumption at the expense of performance, as the system clock frequency does not need to be significantly greater than the sampling frequency in many cases [15]. The power consumption of state-of-the-art ULP SoCs varies greatly with the application, but it is typically on the order of 1 to 10s of μ Ws [9–14].

Table 1.2. Sampling frequency requirements for common IoT sensing applications.

Application	Signal	Sample Frequency	Source
Healthcare			
Heart rate variability	Electrocardiogram (ECG)	100 – 250 Hz	[16]
Blood oxygenation	Photoplethysmography (PPG)	30 – 670 Hz	[17]
Activity recognition	Acceleration	50 – 200 Hz	[18, 19]
Knee joint health	Skin temperature	1 Hz	[20]
	Electrical bio-impedance	22 mHz	
Perspiration		17 mHz	[21]
Environment			
Soil temperature	Temperature	1 mHz	[22]
Soil humidity	Humidity	1 mHz	[22]
Landslide early detection	Vibration	5 – 200 Hz	[23]
Infrastructure			
Concrete health	Resistivity (moisture)	100 mHz	[24]
	Temperature		
Building occupancy	Sound	10 Hz	[25]
Commercial			
Fish spoilage detection	Impedance	4 mHz	[26]

1.2 Thesis

The lifetime and reliability of self-powered IoT SoCs operating in energy-scarce environments will be improved by reducing their power consumption. Enabling self-powered operation for more applications and energy harvesting environments will help address the scaling challenge of battery recharge and replacement hindering the growth of the IoT. This dissertation presents digital circuits, digital circuit design techniques, and system-level strategies for reducing the power consumption of wireless Transmitter (TX) and digital processing circuits in ULP SoCs. This dissertation also presents two ULP SoCs, their applications, and integration into IoT sensing systems. The creation of self-powered systems that demonstrate commercially relevant

IoT applications will accelerate the proliferation of IoT sensing devices by serving as proof of technology readiness.

1.2.1 Reducing Wireless Transmitter Power

Wireless TX circuits often dominate the power consumption of ULP SoCs, and consume significantly more power than other components. The wireless TX in [10] consumes 65% of the 6.45 μ W SoC power budget for an application that intermittently samples and transmits acceleration data from an off-chip sensor. Significant research has been done on the design of wireless TX circuits to reduce their power [JB6,JB9], but techniques outside of the TX design itself must also be considered to fully minimize the TX contribution to SoC power. Typically, TXs only consume high power when actively transmitting data, therefore, reducing the amount of data that must be transmitted using system-level strategies can significantly reduce TX and overall SoC power. This approach is agnostic to the TX design itself, and thus widely applicable to wireless ULP SoCs.

1.2.2 Reducing Digital Circuit Power

Digital circuits are essential for system control and data processing in ULP SoCs, and can be a significant source of power consumption in many applications. Digital circuits consume 36% of the overall power budget for the SoC in [14], which is sampling ECG data at 512 Hz and processing it to extract the heart rate. Static Complementary Metal-Oxide-Semiconductor (CMOS) logic has been the dominant method for implementing digital circuits since its inception in the 1960s [27], and a tremendous amount of research has been done to improve its performance and energy efficiency over the years. High performance isn't required for low-frequency IoT sensing applications though, so frequency-dependant dynamic power is less significant, and static leakage power is dominant in these circuits. This leads to static CMOS circuits such as microprocessor cores being limited by leakage to 10s to 100s of nWs during active-mode operation [11,28,29]. Exploring alternative methods for implementing digital circuits (and design

techniques for those circuits) that can operate below the leakage floor of conventional static CMOS will expand the number applications and operating environments reachable with ULP SoCs.

1.2.3 Accelerating Proliferation of Self-Powered IoT Sensing Devices

Although commercial products exist that use self-powered IoT sensing devices [30–32], they're not prolific or widely adopted yet. This is partly because the value of self-powered operation is not widely understood yet, as it only manifests in very large-scale deployments. It is also partly because reliable self-powered operation in uncertain energy harvesting conditions is challenging, and many of the technologies that enable it are relatively new. Therefore, the creation of self-powered and ULP systems that integrate state-of-the-art components and demonstrate commercially relevant applications will help accelerate the adoption of IoT sensing devices in the commercial space by acting as a proof of technology readiness.

1.3 Contributions

1.3.1 Ultra-Low Power Sensor Data Compression

Chapter 2 presents the design of a lossless sensor data compression accelerator [JB2] for power reduction in wireless IoT devices. This design reduces overall SoC power consumption through the use of on-chip sensor data compression, which reduces the amount of data that must be transmitted, and thus the TX contribution to system power. Although compression does incur some processing power overhead, it results in system-level power reduction because the savings in transmission power is greater than the additional power consumed by compression. The accelerator is implemented on a health-monitoring SoC [JB3,JB7] fabricated in 130 nm CMOS, and closely integrated with the TX interface, realizing a custom data-flow architecture [33] that minimizes the overhead of using compression. The accelerator adds only 4.4 nW processing power overhead to the SoC when operating at 32 kHz and 0.5 V, and achieves an

average Compression Ratio (CR) of 2.39, the lowest power and highest average CR reported to date for lossless on-chip compression of ECG data.¹ For an application that samples ECG data at 360 Hz, the accelerator reduces the required TX duty cycle by 3.7x, the overall system power by 2.9x, and allows the entire system to consume just 2.62 μ W. The overall system power consumption was 7.7 μ W without compression, which is comparable to other state-of-the-art ULP SoCs [9–14], demonstrating that this approach can effectively stack with other aggressive power reduction strategies. This chapter also evaluates the performance of the compression algorithm implemented by the accelerator for biomedical sensor data types (ECG and human acceleration data) for the first time. The algorithm is shown to be effective for most sensor data types with significant temporal correlation between consecutive samples, making it widely applicable to wireless IoT devices.

1.3.2 Ultra-Low Power DVFS RISC-V Microprocessor

Chapter 3 presents the design of a RISC-V microprocessor core implemented using a new Scalable Dynamic Leakage Suppression (SDLS)² logic style [JB10], which enables a continuous power-performance trade-off at power levels below the leakage floor of conventional static CMOS digital circuits. Because energy harvesting sources can fluctuate unpredictably, self-powered SoCs can benefit significantly from the capability to dynamically trade-off power and performance. This allows them to remain operational when harvesting conditions are poor, and to maximize sensor data throughput when harvesting conditions are good. Furthermore, if an SoC can scale its power down to the nW-level, it can operate continuously in even poor energy harvesting conditions with a small energy harvester, and potentially without an energy storage device. This microprocessor core can continuously scale its performance from 6 nW at 11 Hz operating frequency to 140 nW at 8.2 kHz operating frequency at a constant V_{DD} of 0.6 V. Across the supply voltage range, the core is capable of delivering a minimum

¹At the date of publication. More recent work has further improved the state-of-the-art for CR.

²DLS logic is an emerging method for implementing digital circuits that is significantly lower power than conventional static CMOS, but much slower. Chapter 3 introduces DLS logic in more detail.

power of 840 pW, a maximum frequency of 41.5 kHz, and a minimum energy of 13.4 pJ/cycle. This work effectively addresses the poor power-performance scaling granularity of prior DLS microprocessor cores [34, 35], and unlocks a previously inaccessible area of the IoT application space, later shown in Fig. 3.8.

1.3.3 Multi-Threshold DLS Logic

Chapter 4 presents a gate-level Multi-Threshold technique for Dynamic Leakage Suppression logic (MT-DLS) [JB13] that enables further power reduction below the leakage floor of conventional static CMOS digital circuits. MT-DLS is a static technique that enables a power-performance trade-off that is fixed at design time, in contrast to SDLS, which enables a dynamic trade-off that can be adjusted at runtime. An array of 8-bit Multiply-Accumulate (MAC) blocks implemented in 65 nm CMOS demonstrate that MT-DLS yields power savings of up to 6.5x compared to single-threshold DLS logic, and power savings in the 10 Hz–1.4 kHz range across the 0.2 V–1.0 V supply voltage range.

1.3.4 Ultra-Low Power IoT Systems and Applications

Chapter 5 presents two ULP SoCs, their integration into IoT sensing systems, and applications for the purpose of accelerating the commercial adoption and proliferation of IoT devices. First, we present a self-powered SoC [JB3, JB7] and use it to implement several health-monitoring applications, including free fall detection, temperature sensing, and arrhythmia detection using the on-chip Analog Front End (AFE) [JB4]. Then, we integrate the SoC into a vigilant cardiac monitoring system that is powered completely by body heat, and tightly integrate the system into a compact wearable form to enable comfortable long-term use [JB14]. The entire system consumes 65 μ W measured from the energy storage device, and operates in an always-on mode such that no critical Atrial Fibrillation (AFib) events are missed. We also present an ULP wireless wake-up and control system, which is designed to reduce the electricity usage of Miscellaneous Electric Loads (MELs), and includes a fully integrated 802.11ba Wake-Up

Receiver (WuRX) [JB12] and an 85 nW node-controlling SoC [JB11].

Chapter 2

Ultra-Low Power Sensor Data Compression

2.1 Introduction

Data compression is the process of encoding information in a way that uses less bits than the original representation. It is useful because it reduces the memory capacity required to store information, and also the energy required to transmit it, as shown in Fig. 2.1. It is thus very relevant to IoT devices, since the primary function of most IoT devices is to sense, store, and transmit information.

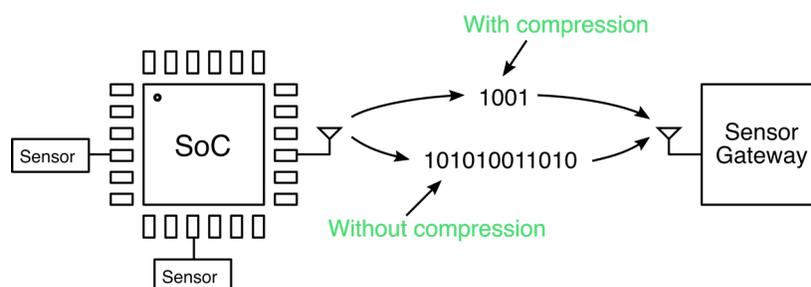


Fig. 2.1. On-chip sensor data compression reduces the amount of data that must be transmitted.

Most IoT devices have very limited memory capacity and energy available to store and transmit information, so they can potentially benefit a lot from data compression. However, they also typically have very limited computational resources and energy available to compress the data. This is exacerbated by the fact that most data compression algorithms are designed for comparatively high performance desktop or mobile computers, which are either powered by the grid or large batteries. Most data compression algorithms are therefore optimized to reduce the size of the data as much as possible, without regards towards computational complexity

The content in this chapter is based on [JB2]. For a list of individual contributions see section 6.3.

or energy efficiency. For IoT devices though, if the energy required to compress the data is greater than the energy saved by storing or transmitting less data, then data compression is not useful and actually detrimental.

There are two types of data compression algorithms: lossless and lossy. Lossless algorithms preserve all of the data, and work by identifying and eliminating statistical redundancy. When a lossless algorithm is reversed, the output data is identical to the input data that was compressed. Lossy algorithms also identify and eliminate statistical redundancy, but they do not preserve all of the data. When a lossy algorithm is reversed, the output data is not identical to the input data that was compressed, as data that is considered not necessary or less important is removed. The performance of compression algorithms is measured by the CR, which is defined as the ratio between the number of bits in the original representation and the number of bits in the compressed representation.

$$\text{CR} = \frac{\text{original number of bits}}{\text{compressed number of bits}} \quad (2)$$

Lossy algorithms typically have much higher CRs than lossless algorithms, but they must be more carefully applied to ensure important information is not lost. Lossy algorithms are often designed and highly customized for a single data type [36]. Lossless algorithms are thus much more widely applicable to multiple data types than lossy algorithms, since all of the information is always preserved. For both lossless and lossy algorithms though, the CR can vary significantly with the characteristics of the data being compressed, even within a particular data type. Thus when assessing the performance of a compression algorithm, it is important to use a large data set that has a realistic amount of variance, similar to what would be encountered in real applications. Statistical measures of the CR such as the mean, variance, minimum, and maximum should be considered. It is also important to use standardized data sets that are publicly available, so that a fair comparison can be done between compression algorithms.

2.2 Motivation

Most IoT devices rely on wireless TXs to communicate with the outside world, as wired communication significantly increases the difficulty and cost of deployment in many environments. Significant research has been done on the design of wireless TX circuits to reduce their power consumption [JB6,JB9]. However, wireless TXs still consume relatively high power compared to other components in ULP SoCs. For the SoC in [10], which performs an application reading and transmitting acceleration data from an off-chip sensor, the TX consumes $4.18 \mu\text{W}$ of the total $6.45 \mu\text{W}$ system power. Therefore, anything that can be done to reduce the TX power will significantly reduce the total system power. Techniques outside of the design of the TX itself must also be considered to fully minimize its power consumption. The benefits of system-level techniques can stack with improvements in the the TX design itself, yielding a compounded effect.

One common system-level technique to reduce TX power is aggressive duty-cycling, but that can severely limit sensor data sampling rates, prohibitively so for many applications. Another common technique is to perform most of the sensor data processing on-chip and only transmit important results. The ULP SoC in [37] collects ECG data and processes it on-chip, only transmitting the data when a critical Atrial Fibrillation (AFib) event occurs. In a sense, this can be considered an extreme form of lossy data compression, and this is a good strategy for the most energy-constrained systems. However, the actual sensor data itself is often more valuable than just the important results for many applications. This is because the processing capabilities of computers connected to the power grid are orders of magnitude greater than that of edge devices such as self-powered SoCs, which typically means more valuable information can be extracted from the data. In addition, having large amounts of unprocessed sensor data available allows for the algorithms that process the data to be improved over time. Many machine learning classifiers, for example, require large amounts of training data to produce good results [38].

On-chip sensor data compression is therefore an ideal strategy for reducing the power of wireless TXs, because it reduces the required TX duty cycle, but still maintains the sensor data sampling rate and allows the unprocessed sensor data to be transmitted. It is also agnostic to the actual TX design itself, and is thus widely applicable to IoT devices. Lossless algorithms are preferable for IoT devices that collect data from multiple types of sensors, since lossy algorithms are typically not applicable to more than one data type. Lossless algorithms are also better for applications like health monitoring, where important diagnostic information could be lost if a lossy algorithm is not applied carefully. This dissertation primarily focuses on lossless data compression for these reasons.

2.3 Prior Art

Several important principles must be considered when evaluating a compression algorithm used for the purpose of minimizing power consumption in ULP SoCs. First, the algorithm must be computationally efficient to minimize the processing overhead of compression. Second, the algorithm must maximize the CR to minimize the amount of data that must be transmitted or stored. And third, the algorithm must account for the specific sensor data, since the CR of a compression algorithm is highly dependent on the characteristics of the data being compressed.

There are multiple fundamentally different approaches to lossless data compression. Dictionary based approaches are used by many popular compression algorithms, and work by searching for matches between the data to be compressed and a set of data contained in a data structure maintained by the encoder. This data structure is referred to as the dictionary, and when a match is found, a reference to the position of the data in the dictionary is substituted in place of the data. One example is the Lempel-Ziv-Welch (LZW) algorithm used in the popular image format Graphics Interchange Format (GIF) [39]. Sensor Lempel-Ziv-Welch (S-LZW) is a variant of LZW that was proposed in 2006 specifically for resource-constrained sensor nodes [40], which adjusts the dictionary size, the size of the data blocks being compressed, and the protocol followed when the dictionary is full to make the algorithm implementation

more energy efficient. S-LZW achieved modest CRs of up 3.5 when applied to a range of environmental sensor data sets.

Dictionary based approaches are inherently resource intensive from a hardware perspective though, at least in the context of ULP SoCs. The minimum dictionary size of 512 entries explored in [40] still requires 2618 bytes of Random-Access Memory (RAM) and 1262 bytes of Read-Only Memory (ROM) to implement. In addition, the operation of comparing the data being compressed with entries in the dictionary to search for a match requires a large number of memory transactions. Fortunately, there are alternatives to dictionary based approaches that only require simple integer-based operations. A large number of such algorithms have been proposed in recent years [41–45], and they all follow a similar formula: a linear prediction step followed by an entropy encoding step. The Lossless Entropy Compression (LEC) algorithm [41] was proposed in 2008, and is one of the first and most well known of such algorithms. In LEC, the prediction step is implemented as a simple delay stage, meaning the prediction is set as the previous sample. The difference between the current and predicted sample (previous sample) is computed, as the differences are much more amenable to entropy encoding than the actual samples themselves. In the entropy encoding step, the differences are appended with prefix codes that represent the size of the differences, thus eliminating the need for framing and allowing the outputs to be variable length.

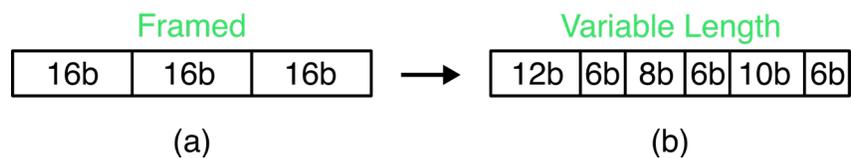


Fig. 2.2. (a) Data framed into 16b words, and (b) variable length data, without the need for framing.

Prefix codes have the property that no shorter codeword may be the prefix of a longer codeword, so they are uniquely decipherable, and thus allow the output to be decoded. Thus, the variable length outputs can be packed together in a continuous bit stream, which is the fundamental reason why compression is achieved in the entropy encoding step.

The LEC algorithm was tested with several environmental sensor data sets for temperature

Prefix Codes
 00, 010, 011, 100, 101,
 110, 1110, 11110, 111110, \longrightarrow **Uniquely Decipherable**
 1111110, etc.

Fig. 2.3. Example binary codes with the prefix property, where no shorter code is the prefix of a longer code.

and relative humidity, and achieved CRs of approximately 3. Using software to execute the algorithm on a Atmel AVR microcontroller, the compression of one input sample requires an average of 355 instructions, 42% of the average number of instructions required by S-LZW. If implemented in hardware though, LEC has the potential to be even more efficient compared to S-LZW, since it only requires integer-based operations and does not require significant memory or memory transactions.

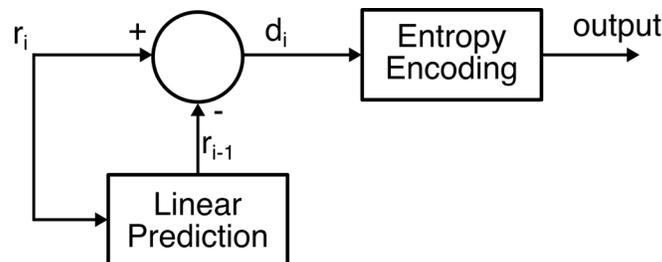


Fig. 2.4. Block diagram of the LEC compression algorithm. The input is represented as r_i , and the difference between the current input and the predicted input is represented as d_i .

The algorithm presented in [42] in 2011 is similar to LEC, but the entropy encoding is performed in a different way using Golomb-Rice coding. The algorithm was tested on several types of biomedical data, and produced a CR of 2.38 on ECG data from the MIT-BIH Arrhythmia Database [46]. However, it is not specified if this CR is an average for all the records in the database, or if it is just for a particular record, or even a small section of a particular record. This algorithm was also implemented as a custom hardware accelerator in a 65 nm CMOS process, with an approximate 54k gate count. The simulated power consumption was 170 μ W when operating at 1 V and 24 MHz.

The algorithm presented in [43] in 2013 (and later in [47]) is also similar to LEC, but explores using higher order linear predictors and does the entropy encoding step in a slightly different

way. The average CR achieved for the entire MIT-BIH Arrhythmia Database was 2.25. The algorithm was implemented as a hardware accelerator in a 0.35 μm CMOS process, with a 2.26k gate count, and the measured power consumption was 2.14 μW at 2.4 V and 32 kHz.

In 2014, the authors of the LEC algorithm proposed several new variants of LEC that add an additional adaptation step to make the algorithm more robust and improve its performance [48]. The new adaptation step involves rotating the prefix code tables used in the entropy encoding step to try to ensure that the shortest prefix codes are used to encode the most frequently occurring differences. This improves the algorithms performance for cases where the mean of the differences distribution is not centered on zero, improving the robustness to different input data characteristics. It is also computationally simple to implement, since the adaption step only requires rotating the prefix code tables. Two methods for managing the prefix code tables were proposed: a Greedy-Adaptive Lossless Entropy Compression (GA-LEC) approach and a Frequency-Adaptive Lossless Entropy Compression (FA-LEC) approach. In addition, two further variants called GAS-LEC and FAS-LEC were proposed that split the prefix encoding tables into two independently rotating data structures, and were further shown to improve the CR. However, the performance of these algorithms were only assessed for several environmental sensor data sets for air temperature and relative humidity, which are phenomena that can change relatively slowly in time. Therefore it is unclear if such favorable results will occur with other types of sensor data that may have less temporal correlation between consecutive samples. Additionally, these algorithms were not implemented in hardware, so the energy efficiency and overall system-level impact for IoT devices remains unknown.

2.4 A 4.4 nW Lossless Sensor Data Compression Accelerator for 2.9x System Power Reduction in Wireless Body Sensors

2.4.1 Contributions

This section extends the state-of-the-art by first evaluating the performance of the data compression algorithm GAS-LEC on biomedical sensor data, specifically ECG and acceleration data obtained from humans performing common activities.¹ Second, the algorithm is implemented as a custom hardware accelerator on a health monitoring SoC [JB3,JB7] fabricated in a 130 nm CMOS process. The accelerator is closely integrated with the wireless TX interface, realizing a custom data-flow architecture [33] that minimizes its contribution to system power, and reduces the software overhead required to use compression. The accelerator adds only 4.4 nW processing power overhead to the SoC when operating at 32 kHz and 0.5 V, and achieves an average CR of 2.39, the lowest power and highest average CR reported to date for lossless on-chip compression of ECG data.² For an application that samples ECG data at 360 Hz, the accelerator reduces the required TX duty cycle by 3.7x, reduces the system power by 2.9x, and allows the entire system to consume just 2.62 μ W. The overall system power consumption was 7.7 μ W without compression, which is comparable to other state-of-the-art ULP SoCs [9–14], demonstrating that this approach can effectively stack with other aggressive power reduction strategies.

2.4.2 Algorithm Definition

The steps of the GAS-LEC algorithm are formally defined here, and shown in Fig. 2.5.

First, the linear prediction step is performed, and the difference d_i between the current sample

¹We chose GAS-LEC over FAS-LEC because the adaptation method is significantly less complex, and therefore lower power to implement in hardware.

²At the date of publication. More recent work has further improved the state-of-the-art for CR.

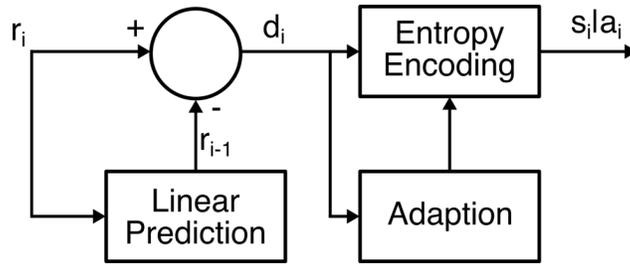


Fig. 2.5. Block diagram of the GAS-LEC compression algorithm.

r_i and the previous sample r_{i-1} is computed.

$$d_i = r_i - r_{i-1} \quad (3)$$

Then in the entropy encoding step, n_i , the minimum number of bits needed to represent d_i is computed.

$$n_i = \begin{cases} 0, & \text{if } d_i = 0 \\ \log_2(|d_i|) + 1 & \text{otherwise} \end{cases} \quad (4)$$

Once n_i has been determined, it is then necessary to calculate a_i , which is simply the n_i lower bits of d_i if d_i is positive. If d_i is negative, then the two's complement of d_i is used. If d_i is 0, a_i is not needed.

$$a_i = \begin{cases} \text{not needed} & \text{if } d_i = 0 \\ (d_i)_{n_i}^L & \text{if } d_i > 0 \\ (d_i - 1)_{n_i}^L & \text{if } d_i < 0 \end{cases} \quad (5)$$

The prefix code s_i can then be determined from n_i . Each n_i has a corresponding s_i according to Tables 2.1 and 2.2.

The final compressed output can then be computed, and is defined as s_i concatenated with a_i , or $s_i | a_i$. Next, the adaptation step is performed, which tries to ensure that the shortest prefix codes are used to encode the most frequently occurring differences. In the greedy adaptive approach used by GAS-LEC, the tables are rotated after each d_i is encoded to ensure that the

Table 2.1. Lower Prefix Table

n_i	s_i	d_i
0	00	0
1	010	-1,+1
2	011	-3,-2,+2,+3
3	1110	-7,...,-4,+4,...,+7
4	11110	-15,...,-8,+8,...,+15
5	110	-31,...,-16,+16,...,+31
6	100	-63,...,-32,+32,...,+63
7	101	-127,...,-64,+64,...,+127

Table 2.2. Upper Prefix Table

n_i	s_i	d_i
8	111110	-255,...,-128,+128,...,+255
9	1111110	-511,...,-256,+256,...,+511
10	11111110	-1023,...,-512,+512,...,+1023
11	11111110	-2047,...,-1024,+1024,...,+2047
12	1111110	-4095,...,-2048,+2048,...,+4095

center of each table is on the group belonging to the previous sample. Note that the center of each prefix table is defined as the location of the shortest prefix code s_i , and that the length of the prefix codes increases outwards from the center. This ensures that differences closer to the center of the table are assigned the shorter prefix codes, which means that differences that are closer to the previous differences are encoded with shorter prefix codes.

2.4.3 Algorithm Performance

We implemented the algorithm in software and evaluated its performance using the MIT-BIH Arrhythmia Database [46] for ECG and the UCI Machine Learning Repository [49] for acceleration. We first evaluated the effectiveness of the linear prediction step by looking at the range of the differences compared to that of the actual samples, shown in Fig. 2.6. We also computed the mean and variance of the resulting difference distributions for one ECG data set and several acceleration data sets for different activities shown in Table 2.3. The mean is not shown in the table, but was reduced to approximately zero for all of the difference distributions. We then executed the complete algorithm on numerous ECG and accelerometer sets, shown in

Fig. 2.7 and 2.8. Because we achieved average compression ratios of 2.4 and 2.0 for ECG and acceleration, respectively, we conclude that the algorithm effectively compresses such biomedical sensor data.

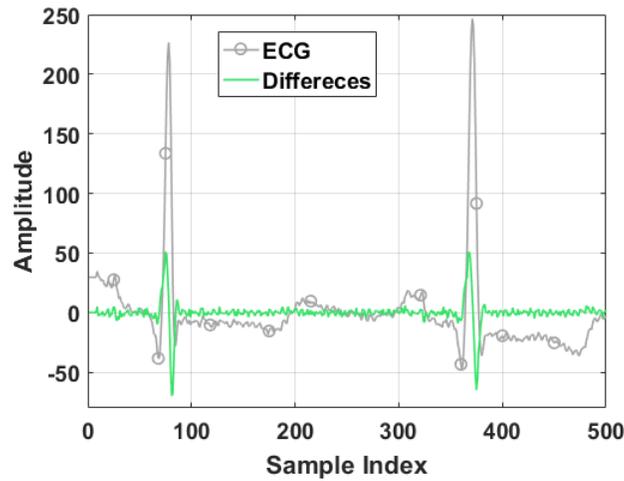


Fig. 2.6. The range and variability of the differences is much less than that of the actual ECG signal. The data is from MIT-BIH record 100 and the DC offset was removed from the ECG waveform to better compare the variability.

Table 2.3. Statistics showing the effectiveness of the linear prediction step on ECG and acceleration data.

Data Set	Variance of Samples	Variance of Difference	Improvement
ECG record 100	1388.8	114.3	12.2x
Working at computer	730.4	109.2	6.7x
Talking while standing	738.3	49	15.1x
Standing	1068.3	103.1	10.4x
Walking	4249.2	1062.0	4.0x
Walking and talking	5464.6	1046.4	5.2x
Going up/down stairs	6766.8	861.9	7.9x
Standing, walking, stairs	6846.8	371.0	18.5x

We also implemented several other compression algorithms in software to compare their performance against GAS-LEC for biomedical sensor data. The results are shown in Table 2.4, and clearly show that GAS-LEC outperforms the other algorithms across the board for acceleration and ECG biomedical sensor data.

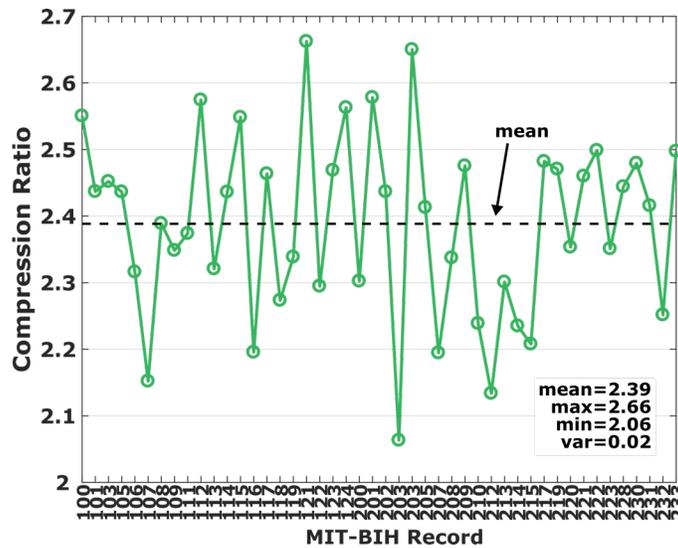


Fig. 2.7. The proposed algorithm is effective at compressing 46 ECG records from the MIT-BIH database. Modified lead II (MLII) data was used.

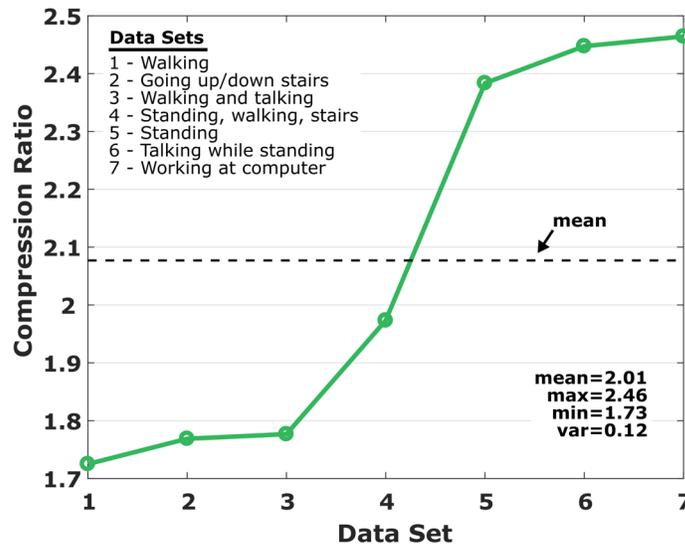


Fig. 2.8. The proposed algorithm is effective at compressing acceleration data for various activities of daily living from the UCI repository. Data from participant 1 was used.

2.4.4 Implementation and Measured Results

The accelerator was carefully designed to minimize the processing power increase and user overhead. The accelerator was integrated into the transmitter interface on a health-monitoring SoC [JB3], shown in Fig. 2.9, and can be bypassed or enabled with a single configuration bit. This allows the accelerator to be used without any additional software instructions, except the

Table 2.4. CR comparison for state-of-the-art compression algorithms for ECG and various acceleration data sets.

Algorithm	ECG	Walking	Standing	Working at Computer	Standing, Walking, Stairs
LZW (512B dict.)	1.13	1.16	1.88	2.10	1.28
LZW (2048B dict.)	1.20	1.34	1.96	2.37	1.31
LEC	2.63	2.14	3.06	3.12	2.49
GA-LEC	1.82	1.62	2.02	2.04	1.76
GAS-LEC	2.74	2.29	3.22	3.28	2.62

one instruction required to enable it. Without this feature, the overhead of moving data in and out of the accelerator could significantly increase the system power, since more clock cycles are required to move the data and more memory capacity is required to store the additional instructions. The accelerator relies mostly on combinational logic, and processes inputs in two clock cycles, adding minimal latency to the system. It also has a special operating mode that allows it to effectively compress multiple sensor data streams concurrently, such as 3-axis acceleration data.

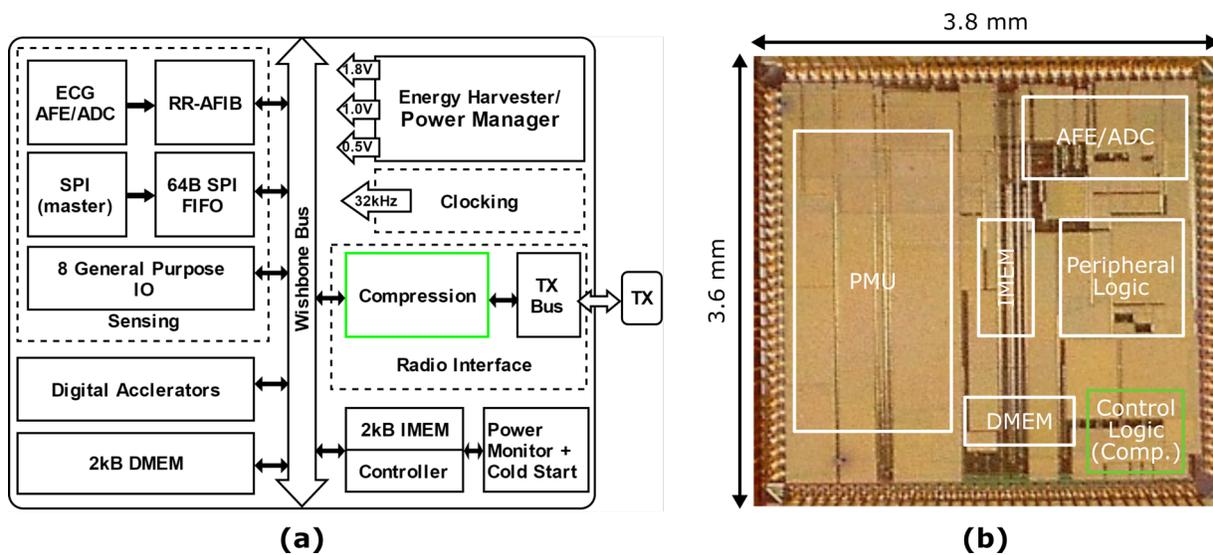


Fig. 2.9. (a) Block diagram of the SoC, showing the compression accelerator is on the TX interface data path. (b) Annotated die photograph of the SoC, showing the compression accelerator (part of control logic).

To evaluate the system-level benefits of compression, we measured the system power consumption with and without compression while transmitting ECG and acceleration sensor data. The system uses an off-chip Frequency Shift Keying (FSK) modulated TX that works at 2.4 GHz and is controlled by the SoC through the TX interface. We measured the active

power of the TX to be 1.05 mW and the sleep power to be 160 nW. For ECG, we used data from MIT-BIH record 100 and maintained a data rate of 360 Hz. For acceleration, we used data from UCI participant one for the activity of standing, walking, and going up/down stairs and maintained a data rate of 52 Hz (per axis with 3-axis data, 156 Hz overall). In software, a CR of 2.5 and 2.0 was achieved for the ECG and acceleration data sets, respectively. In hardware, the benefits of compression are compounded due to the system architecture being 16b and the sensor resolutions being 11b and 12b for ECG and acceleration respectively. This is a common occurrence in many systems, and has the effect of increasing the uncompressed size of each sample to the system bit-width unless additional data packing techniques are applied. Therefore, the compression accelerator achieves a TX duty cycle reduction of 3.7x and 2.6x for ECG and acceleration, respectively. This reduces the power of the TX by 3.5x and 2.4x for ECG and acceleration, respectively, as shown in Fig. 2.10. Enabling compression increases the processing power by only 4.4 nW at 0.5 V and 32 kHz for ECG, and compression reduces switching activity in the SoC TX interface as well, thus the system power is reduced without even including the savings in the TX itself. The total system power is reduced by 2.9x and 2.0x, allowing the system to consume only 2.62 μ W and 1.87 μ W while transmitting ECG and acceleration data respectively. Due to area and pad constraints, the compression accelerator was synthesized with and shares a supply voltage with other digital blocks. Therefore, we were unable to obtain the gate count or measure the static part of the power. As shown in Table 2.5, this implementation achieves the lowest power reported to date and provides significant system-level power savings with minimal overhead.

2.4.5 Conclusion

The contributions presented in this section demonstrate that the lossless compression algorithm GAS-LEC is effective at compressing both ECG and human acceleration biomedical sensor data. They also show that it can be implemented with a very low processing power overhead, and that it can significantly reduce the wireless TX power, and thus system power if implemented carefully.

Ultimately, the system power reduction in ULP SoCs enabled by sensor data compression will improve the capabilities of IoT devices to operate reliably from energy harvesting in uncertain environments, and help further the proliferation of IoT devices.

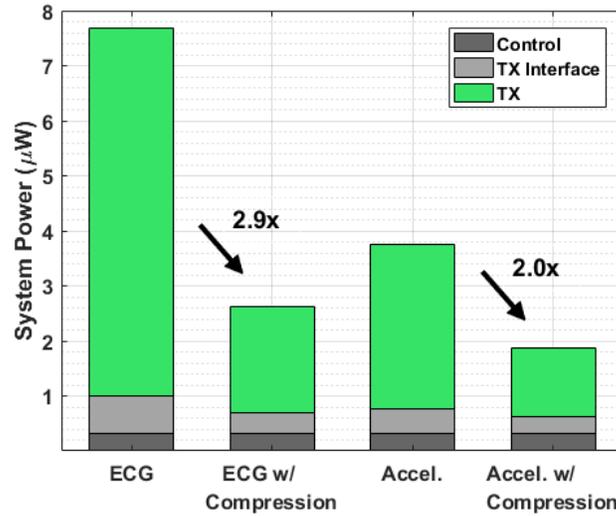


Fig. 2.10. System-level results, showing the power savings of compression. The SoC power was measured and the TX power was calculated based on the duty cycle and previous measured results for active and sleep power.

Table 2.5. Comparison with state-of-the-art on-chip lossless data compression accelerators.

	This Work [JB2]	TCE'11 [42]	JSSC'14 [47]
Technology	130 nm	65 nm	0.35 µm
Supply Voltage (V)	0.5	1.0	2.4
Frequency	32 kHz	24 MHz	32 kHz
CR ECG (MIT-BIH)	2.39	2.38	2.25
CR Accel. (UCI)	2.01	–	–
Data Types Supported	ECG/Accel. Most sensors	ECG	ECG/EEG DOT
Power (ECG)	4.40 nW	170 µW	2.14 µW
System Power	2.62 µW ECG 1.87 µW Accel.	–	–
System Power Reduction	2.9x ECG 2.0x Accel.	–	–

Chapter 3

Ultra-Low Power DVFS RISC-V Microprocessor

3.1 Introduction

The conventional static CMOS logic family was first invented in 1963 by engineers at Fairchild Semiconductor [27]. It was originally proposed as a low standby power alternative to other logic families of the time, as the power consumption when no switching operation is occurring is due only to leakage current. This is in contrast to other competing logic families of the time, such as NMOS only logic, which has significant active current in the steady state. Today, almost 60 years later, static CMOS is the most widely used logic family for implementing digital circuits.

The total power consumption of a static CMOS digital circuit can be represented as the sum of the dynamic and static components, shown below.¹

$$P_{\text{total}} = P_{\text{dynamic}} + P_{\text{static}} = f_{\text{clk}} \alpha C_{\text{eff}} V_{\text{DD}}^2 + I_{\text{leak}} V_{\text{DD}} \quad (6)$$

The dynamic part refers to the power dissipated when switching activity occurs, and is proportional to the clock frequency f_{clk} , the activity factor α , the effective switched capacitance C_{eff} , and the square of the supply voltage V_{DD} . The static part refers to the power dissipated when no switching activity occurs, and is proportional to the leakage current I_{leak} , and the supply voltage V_{DD} . In high-performance high-frequency (GHz–MHz) circuits, the dynamic power typically dominates and makes the static power inconsequential. However, the advent of the IoT has led to the emergence of many applications for low-frequency (kHz–Hz) circuits, in

The content in this chapter is based on [JB10]. For a list of individual contributions see section 6.3.

¹This equation neglects short-circuit current. Power due to short-circuit current can be considered part of dynamic power consumption, and is typically less than 10% of other sources of dynamic power [50].

which the static power is significant compared to the dynamic power. Technology scaling also contributes to making static power more significant, as shrinking device dimensions lead to reduced parasitics, and thus less dynamic power [51].

There are many techniques for reducing power consumption in static CMOS digital circuits. Supply voltage scaling is one of the most powerful, since dynamic power is quadratically proportional to voltage. Frequency scaling is also powerful, since dynamic power is linearly proportional to clock frequency. The maximum frequency a circuit can achieve depends on the supply voltage, and lower voltages lead to lower maximum frequencies. Supply voltage and frequency scaling are thus often used together, in a technique known as Dynamic Voltage and Frequency Scaling (DVFS), where both are dynamically adjusted at runtime to trade-off power and performance in response to the required work load of the circuit. In this context, the work load refers abstractly to the number of operations a circuit must complete in a given period of time.

The work load is sparse for many IoT applications, meaning the circuits performing them spend most of their time being idle and not performing useful operations [15]. In these cases, power can be minimized by introducing multiple modes of operation, such as an active mode and a standby mode. The overall power consumption thus depends on the time spent in each mode and the power consumption of each mode.

$$P_{\text{total}} = P_{\text{active}}T_{\text{active}} + P_{\text{standby}}T_{\text{standby}} \quad (7)$$

The active time is defined such that $0 \leq T_{\text{active}} \leq 1$, and the standby time is defined as $1 - T_{\text{active}}$. In cases where the application permits $T_{\text{standby}} \gg T_{\text{active}}$, the overall power is determined mostly by the standby power, so it is important to minimize. The dynamic power can often be reduced or eliminated in the standby mode by clock gating, which disconnects the circuit from the clock source. Clock gates are typically realized using a combinational gate with a latch to prevent clock glitching, as shown in Fig. 3.1, and are thus low overhead to implement. If clock gating is effectively utilized, the standby power is set by the static power.

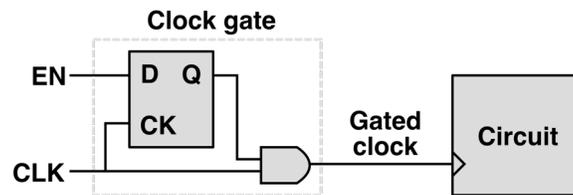


Fig. 3.1. Block diagram of a typical clock gate cell.

A number of techniques for reducing static power are available, including the use of high threshold voltage devices, and cutoff structures. The threshold voltage of a transistor (V_{TH}) refers to the value of the gate-source voltage V_{GS} where strong inversion occurs, and is a function of material constants set during the manufacturing process, and the source-bulk voltage V_{SB} [50]. Most commercial technologies include devices with different V_{TH} , such as low- V_{TH} (L_{VT}), regular- V_{TH} (R_{VT}), and high- V_{TH} (H_{VT}) devices. H_{VT} devices have significantly larger delays than R_{VT} devices at the same nominal V_{DD} due to reduced overdrive voltage ($V_{GS} - V_{TH}$), but they have exponentially lower I_{OFF} current and thus static power consumption, as shown in Fig. 3.2. I_{OFF} refers to the drain current I_{DS} when V_{GS} is equal to 0 V. In addition to using H_{VT} devices that are fixed at design-time, body biasing can also be used to adjust V_{TH} through V_{SB} and dynamically trade-off delay and leakage current at runtime. Reverse Body Biasing (RBB) has been shown to significantly reduce static power [52], but has significant overheads, as it requires a special triple-well process, and additional voltage regulators to generate the bias voltages.

A cutoff structure can be used to optionally disconnect a circuit from the supply voltage and reduce its leakage current. They're often implemented with a PMOS header or NMOS footer (sleep transistors), as shown in Fig. 3.3, and can be dynamically turned on or off depending on the operating mode. Cutoff structures can be utilized at different levels of the design hierarchy: at the gate level where each standard cell has an integrated sleep transistor, at the local level where each circuit block has a single sleep transistor, or at the global level where multiple blocks share a sleep transistor [53]. Typically sleep transistors utilize H_{VT} devices to maximize leakage current reduction. The magnitude of the leakage current reduction depends on the size

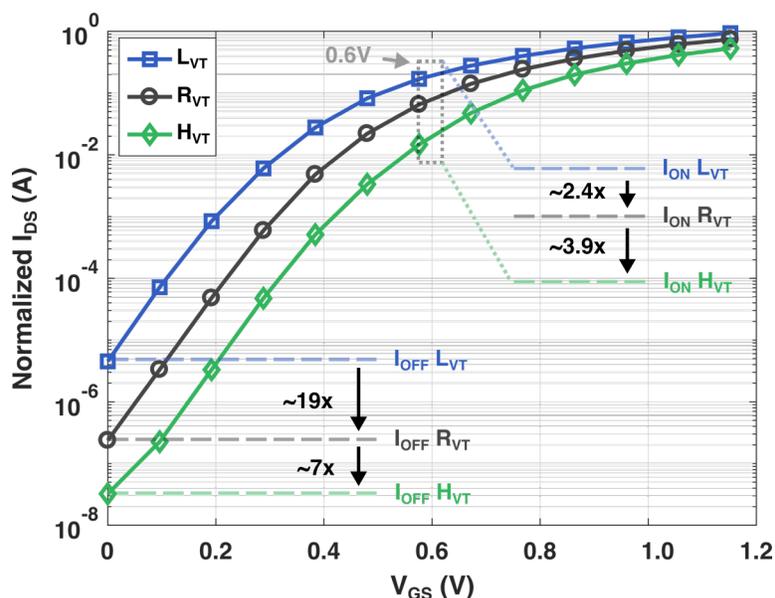


Fig. 3.2. The normalized I_{DS} as a function of V_{GS} for different V_{TH} devices.

of the sleep transistor and its V_{TH} relative to the devices being cutoff. Sleep transistors have the downside of increasing delay when the circuit is not in standby mode though, since they act like a resistor and reduce the effective supply voltage seen by the circuit. The magnitude of the delay increase depends on similar factors as the leakage current reduction. A delay increase of 10% relative to the case without the sleep transistor is a typical rule-of-thumb used for sleep transistor sizing [54]. Cutoff structures also have the downside of being non-retentive, meaning the state of sequential circuits is lost when the standby mode is turned on. This can be addressed by using some form of retention register that uses the global V_{DD} to hold the state instead of the virtual V_{DD} , but this can add significant power, area, and design complexity overhead.

3.2 Motivation

Despite all of the available power reduction techniques, static CMOS circuits are inherently unable to satisfy all of the requirements of the IoT application space. For the IoT application space to be fully realized, the size of IoT devices must be reduced as much as possible to minimize the cost and difficulty of deployment. This leads to smaller energy harvesters and

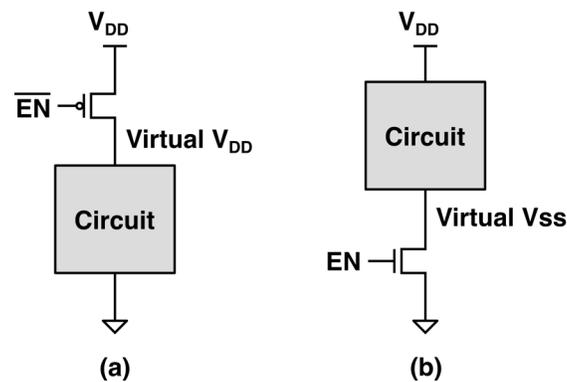


Fig. 3.3. Block diagram of typical cutoff structure implemented using a (a) PMOS header and (b) NMOS footer.

lower power budgets. There is also a strong incentive to ensure reliable operation in as many environments as possible, which necessitates that devices can operate in the worst-case harvesting conditions for extended periods of time. If the state-of-the-art can be improved so that a device can actively operate with nW-level power consumption, many new applications become possible. For example, a $4 \times 4 \text{ mm}^2$ solar cell can support a nW-level device even under extremely low light conditions (11 lx) like moonlight [55].

The minimum power floor for a static CMOS digital circuit is set by the static leakage current. For example, state-of-the-art microcontroller-class processor cores have a minimum power consumption in the range of 10s to 100s of nW when they are in an active state due to leakage [11, 28, 29]. Although their power consumption can be lowered to sub-nW using standby modes with cutoff structures [28], this implies intermittent operation. Such duty-cycled operation requires an energy storage device to act as a buffer, as the active state power is significantly higher than the harvested power. Energy storage devices increase the physical size of the system, and if there is an extended period of time when harvesting conditions remain poor, the stored energy will eventually be depleted and the device will be forced to stay in non-operational standby mode until conditions improve. This interruption in active operation can be prohibitive for some applications, and thus to ensure reliable operation, energy storage devices must be sized very pessimistically to ensure they can sustain the device for the worst-case duration of poor harvesting conditions.

A state-of-the-art commercially available rechargeable thin film battery has a capacity

of 5 μAh at 3.8 V in a 1.70x2.25x0.20 mm^3 package, or a capacity of 50 μAh at 3.8 V in a 5.70x6.10x0.20 mm^3 size [56]. These batteries are similar in size to the die of IoT SoCs [10,11], so they can be integrated without drastically increasing the device size. However, they are only rated for 1,000 recharge cycles at 50% depth-of-discharge at 25 °C, after which their capacity will have degraded to 80%. Thin film batteries are therefore not a viable solution for IoT devices deployed in environments where the harvested power drops below the active system power frequently, as they would need to be replaced often due to degraded capacity.

Super-capacitors are an alternative to thin film batteries that can support significantly more recharge cycles. The thinnest commercially available super-capacitor has a capacitance of 70 mF, a rated voltage of 2.25 V, and is 20x10x0.4 mm^3 [57]. Assuming that all of the stored energy could be extracted without any conversion loss, this could provide 4.92 μW for a 10 hour duration. This particular super-capacitor is rated for 50,000 recharge cycles before a 30% loss of capacitance occurs, so it can support a much longer device lifetime than a thin film battery. However, it is significantly larger than a typical IoT SoC. The IoT SoC presented in [11] is 1.94x1.94 mm^2 , and it is implemented in a 65 nm CMOS technology. This particular super-capacitor would thus increase the system area by 53x.¹ If a more advanced technology node is used, a super-capacitor would increase the system area by an even larger factor.

Regular ceramic capacitors can also be used as an alternative to super-capacitors, as they are available in significantly smaller sizes with reduced capacity to store energy. A standard 0402 surface mount capacitor is just 1.00x0.50x0.35 mm^3 , and is readily available with a capacitance of up to 22 μF at a rated voltage of 4 V [58]. Assuming that all of the stored energy could be extracted without any conversion loss, this could provide 2.93 μW for a 1 minute duration. While this is sufficient to enable duty-cycled operation if the average power consumption stays below the harvested power, if the environmental conditions change for the worse, the system will quickly deplete the stored energy and be forced to stay in a standby mode until conditions improve. And microprocessor cores with sub-nW standby modes are

¹ Super-capacitors exist with smaller areas than the one presented here, but they have similar volumes, so they still significantly increase the system size.

realized only using cutoff structures [28], which implies that they are non-operational in the standby mode, unable to sense or process any information. Transitioning between active and standby modes also often incurs energy and delay penalties, and depending on the system, this can limit the duty cycle and thus average power consumption that can be achieved.

With these limitations in energy storage devices, IoT devices can have a significantly smaller size and longer lifetime if they're able to operate without them. They can also perform applications more reliably if they're able to remain active instead of being forced into a non-operational standby mode when environmental conditions change, and harvested power drops to the nW-level. Ideally, they would also be able to adapt to the changing environmental conditions and trade-off power and performance in a DVFS-like manner, maximizing functionality while keeping the power consumption below the harvested power.

3.3 Prior Art

In 2007, a new ultra-low leakage CMOS logic style was proposed for low-frequency circuits [59]. This logic style was later named Dynamic Leakage Suppression (DLS) logic [34], and it is similar to static CMOS, except the logic gates have an additional NMOS header and PMOS footer transistor. The gates of the DLS header and footer transistors are connected to the output, and are thus self-biasing. Fig. 3.4 shows the DLS inverter logic gate.

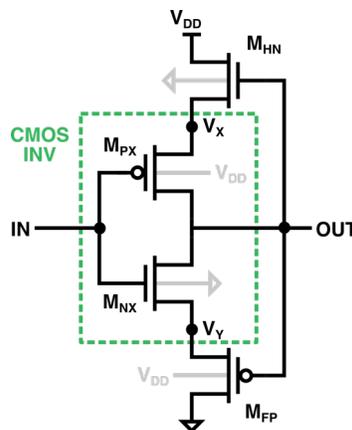


Fig. 3.4. The DLS inverter logic gate.

The leakage current is reduced through the super-cutoff effect, wherein a negative V_{GS} voltage occurs for either the transistors in the pull-up or pull-down network, depending on the output of the gate. Fig. 3.5 shows the DLS inverter logic gate for both steady state conditions. If the output is logic 0, the pull-up network is leaking and there is a voltage drop of V_{DD} across the series connected M_{HN} and M_{PX} transistors. The voltage V_X can thus be solved for by equating the currents between M_{HN} and M_{PX} . Neglecting short-channel effects, if the V_{TH} of the NMOS and PMOS are symmetrical, and the transistors are sized such that the current at threshold is equivalent ($I_{0p}=I_{0n}$), then V_X settles at $V_{DD}/2$. Thus $V_{GS,HN} = -V_{DD}/2$ and $V_{SG,PX} = -V_{DD}/2$. A similar analysis can be done for the opposite case, where the output is logic 1 and the pull-down network is leaking.

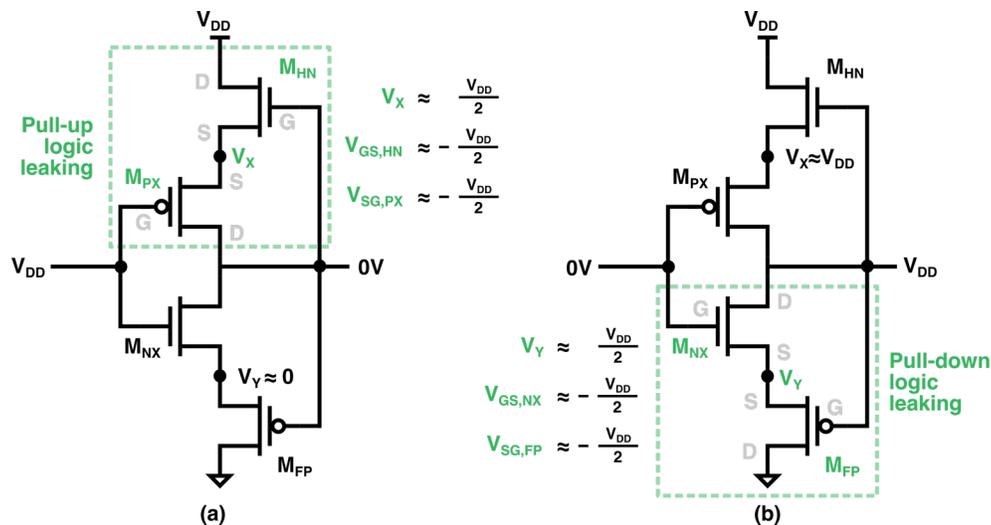


Fig. 3.5. The DLS inverter in steady state, (a) with the pull-up logic leaking, and (b) with the pull-down logic leaking.

The outputs of DLS logic gates exhibit a hysteresis effect, shown in Fig. 3.6, and have sharp transitions between logic levels. For the case of a rising input, the transition occurs when the input voltage exceeds $V_{DD}/2$, and the V_{GS} for transistor M_{NX} increases. This starts to short V_Y with V_{OUT} , which quickly increases the V_{GS} of M_{FP} , leading to the sharp transition. A similar analysis can be done for the opposite case of a falling input transition. The results presented in [59] show that for a 130 nm technology, a DLS inverter has a leakage current of just 2.5 pA at 0.5 V compared to 2.6 nA for a static CMOS inverter. The delay of the DLS

inverter is significantly greater though, at $4.0\ \mu\text{s}$ compared to $0.16\ \text{ns}$ for a static CMOS inverter. The leakage current of a DLS inverter has a minimum at approximately $0.5\ \text{V}$. Although higher voltages lead to reduced subthreshold leakage current due to the V_{GS} of the leaking devices becoming more negative, drain-to-body junction leakage becomes greater than subthreshold leakage around $0.45\ \text{V}$, and limits the achievable minimum leakage current. The supply voltage where minimum leakage current occurs is thus set where the subthreshold and junction leakage are approximately equal. In newer technologies, this isn't the case, as gate leakage becomes increasingly significant due to aggressive oxide thickness scaling.

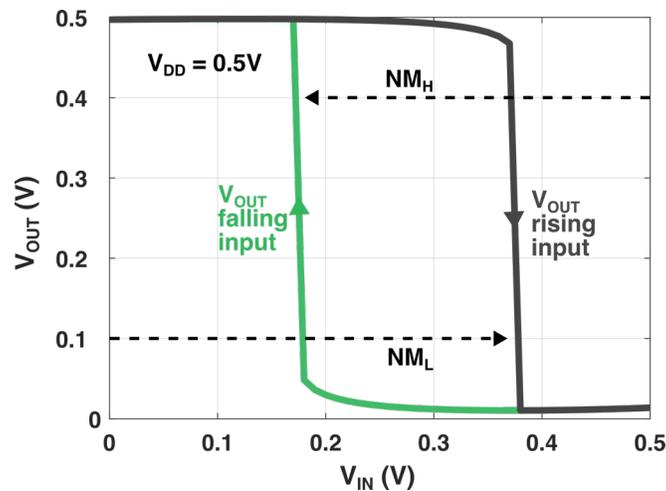


Fig. 3.6. The Voltage Transfer Characteristic (VTC) for a DLS inverter logic gate.

In 2015, a Cortex-M0+ microprocessor core was implemented using DLS logic in a $180\ \text{nm}$ technology and achieved $295\ \text{pW}$ active state power consumption at $550\ \text{mV}$ and $2\ \text{Hz}$ operating frequency [34]. This was the lowest active state power consumption ever realized for a microprocessor core, and the core could be powered with just a $0.09\ \text{mm}^2$ solar cell in dim indoor light ($240\ \text{lx}$). However, the maximum frequency was just $15\ \text{Hz}$, which severely limits the usefulness of the core, especially for the case where energy harvesting conditions are good and more energy is available. In addition, the core wasn't shown performing any actual computations, which indicates that there was a functionality problem with the implementation. The only evidence presented of it being operational was several of the bits toggling on the bus connecting the core with the memory, the minimum necessary to show that the core was

fetching and executing instructions.

In 2018, a new dual-mode variant of DLS logic was presented with two additional bypass transistors in the header/footer that allow the gates to be reconfigured at runtime in either a DLS or static CMOS operating mode [35]. Fig. 3.7 shows the dual-mode DLS inverter logic gate. When the mode bit is logic 0, transistors M_{CN} and M_{CP} are turned on and bypass the M_{HN} and M_{FP} transistors, putting the gates in a static CMOS mode. To further improve performance in this mode, M_{CN} is over-driven by $V_{DD} + \Delta V$, and M_{CP} by $-\Delta V$. When the mode bit is logic 1, transistors M_{CN} and M_{CP} are turned off and the gate operates in DLS mode. A 16-bit MSP430 compatible microprocessor core was implemented with the dual-mode logic gates, and shown to achieve a minimum power of 595 pW (0.45 V, 2 Hz) in the DLS mode, and a maximum frequency of 2.8 MHz (1.1 V) in the static CMOS mode. In the static CMOS mode, the minimum energy was 14 pJ/cycle at 19 kHz, and the minimum power was 118 nW. This is a significant improvement over [34], as it has a similarly low active state power floor, but isn't limited a few Hz in the cases where ample harvested power is available. The poor scaling granularity is limiting though, as 118 nW is a 198x increase from 595 pW. If an energy harvester is minimally sized to provide the system power floor under the worst-case conditions, conditions must substantially improve to support the CMOS mode of operation. Otherwise, duty cycling must be used to achieve an overall power consumption that is an average of the two modes, and that requires an energy storage device, and adds significant overhead to manage to transitions between modes.

3.4 A 6–140-nW 11-Hz–8.2-kHz DVFS RISC-V Microprocessor Using Scalable Dynamic Leakage Suppression Logic

3.4.1 Contributions

This section presents a RISC-V microprocessor core implemented using a novel Scalable Dynamic Leakage Suppression (SDLS) logic style. Together with a custom Adaptive Clock

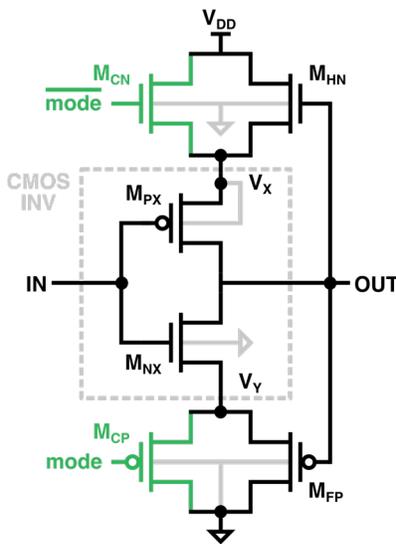


Fig. 3.7. The dual-mode DLS inverter logic gate.

Generator (ACG) and Voltage Scaling Controller (VSC), the core realizes a fully integrated modified DVFS scheme that enables nW-level performance flexibility for self-powered IoT sensor nodes in energy-scarce environments. The core is implemented in a 65 nm low-power CMOS process, and at the nominal V_{DD} of 0.6 V, the core can scale its performance from 6 nW at 11 Hz to 140 nW at 8.2 kHz. Across the supply voltage range, the core realizes a minimum power of 840 pW, a maximum frequency of 41.2 kHz, and a minimum energy of 13.4 pJ/cycle. This work effectively addresses the poor scaling granularity of the dual-mode cells [35], and unlocks a previously inaccessible area of the IoT application space, shown Fig. 3.8. The achievable Hz-kHz frequency range is ideal for many low-frequency IoT applications.

3.4.2 Scalable Dynamic Leakage Suppression Logic

Fig. 3.9 shows the proposed SDLS logic design, which implements a new voltage-scaling approach in which complementary bias voltages V_{CN} and V_{CP} are used to transition the gate across a continuous range between DLS and static CMOS operating regimes, trading-off leakage power and delay. Compared to traditional DVFS, this constant- V_{DD} approach eliminates the need for level shifters and avoids regulator efficiency losses caused by varying supply voltages.

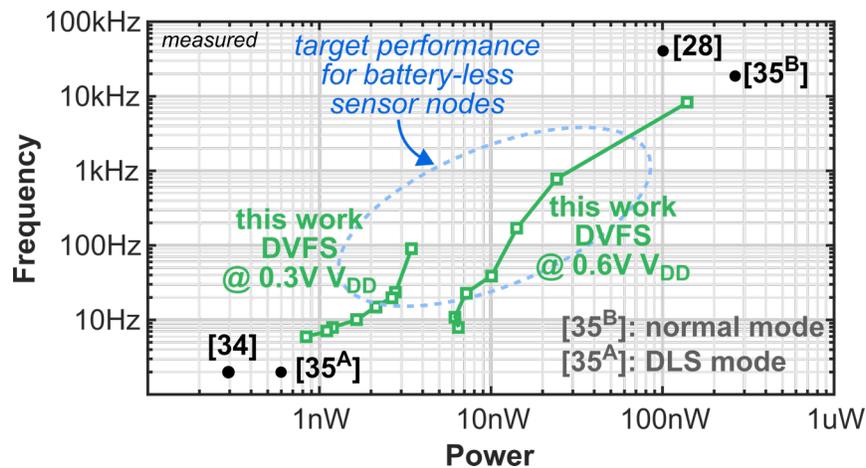


Fig. 3.8. Measured performance with comparison with the state-of-the-art low-power processors.

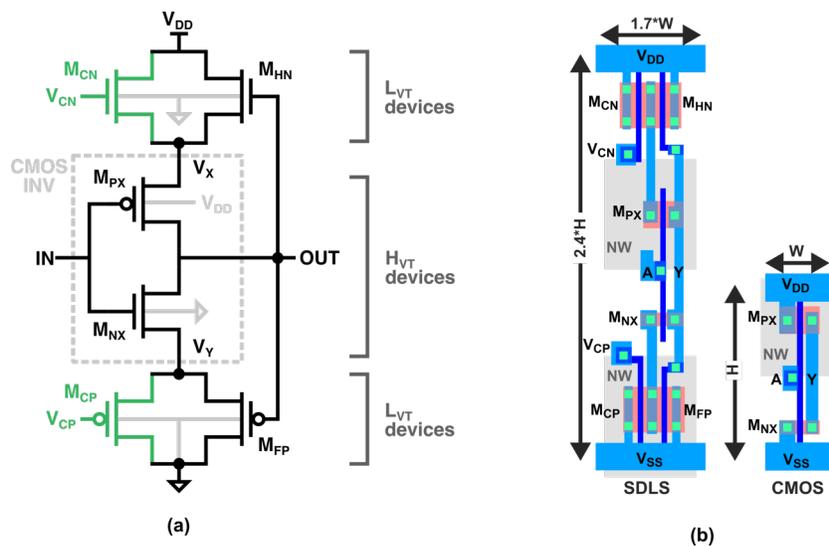


Fig. 3.9. (a) The SDLS inverter logic gate schematic, and (b) the SDLS inverter logic gate layout compared to a standard static CMOS inverter logic gate.

Compared to the dual mode standard cells [35], we also introduce two further design modifications. First, the body of internal PMOS transistors (M_{PX}) are tied to V_{DD} , rather than node V_X , in order to save area by sharing internal n-wells between adjacent cells with negligible performance degradation. Second, L_{VT} devices are used for the external transistors M_{HN} , M_{FP} , M_{CN} , and M_{CP} , allowing increased sensitivity to V_{CN} and V_{CP} for a larger performance tuning range, while simultaneously decreasing the large transistor sized required in [34, 35].

The transient operation of an SDLS inverter is shown in Fig. 3.10, for the case of a falling input transition for three different control voltage values. The control voltage V_C is defined such

that $V_C = V_{CN}$ and $V_{CP} = 1 - V_{CN}$.

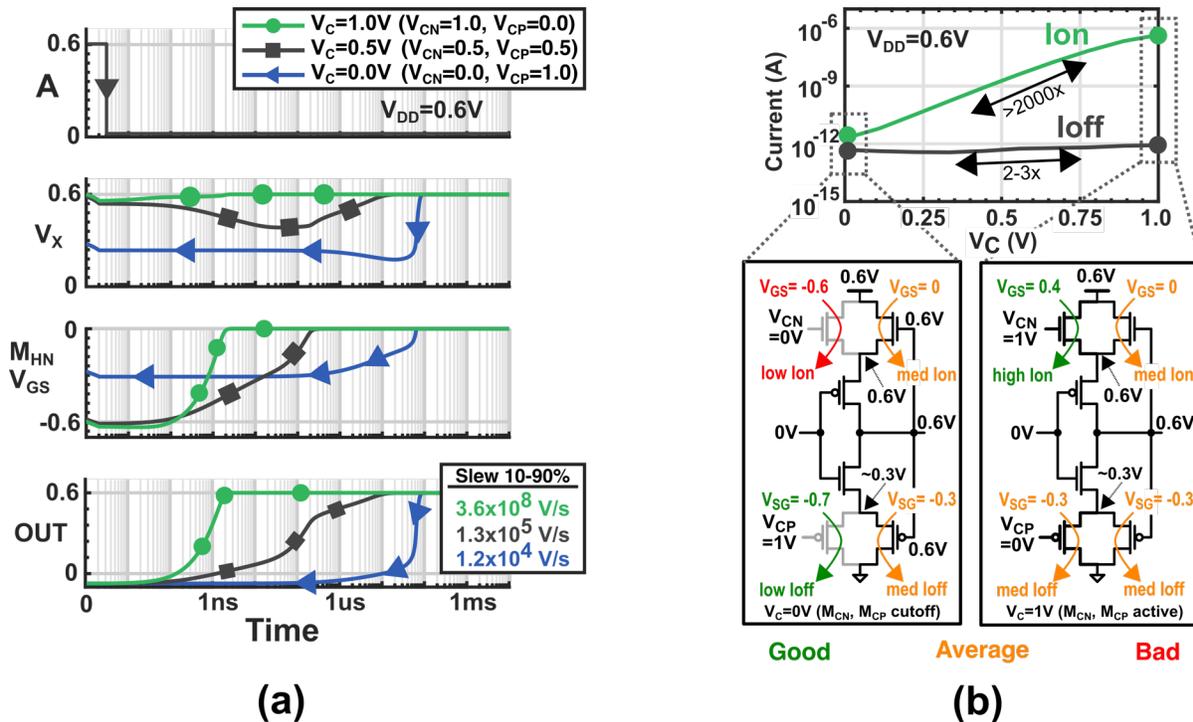


Fig. 3.10. (a) Transient operation of an SDLS inverter logic gate. (b) The bias voltage effect on steady state voltages and on/off current.

In the steady state when the input A is high and V_C is 0 V, the internal node n1 settles to $V_{DD}/2$, and M_{CN} , M_{HN} , and M_{PX} enter super-cutoff mode (negative V_{GS}). When A transitions low, M_{PX} turns on, and V_X converges with the output Y. This creates positive feedback by increasing V_{GS} of M_{HN} , allowing more current to leak through the pull-up network, further charging Y until it has fully transitioned. If V_C is increased, transistor M_{CN} pulls the steady state voltage of V_X higher than $V_{DD}/2$. This weakens the super-cutoff effect in the pull-up network causing increased leakage, and also accelerates the convergence of Y and V_X , allowing a quicker feedback response from M_{HN} that results in a shorter gate delay. In addition to increasing M_{HN} 's feedback response, M_{CN} provides increased on-current throughout the transition that further improves speed. A similar analysis can be done for the opposite case of a rising input transition.

3.4.3 System Architecture

Fig. 3.11 shows the system architecture. It features a 32 bit BottleRocket RISC-V microcontroller-class processor core that implements the RV32IMC instruction set [60]. The core interfaces to a custom 8 kB 6T Static Random-Access Memory (SRAM) macro and an uncore domain that includes system interconnect, Serial Peripheral Interface (SPI) master and General-Purpose Input Output (GPIO) peripherals, and a memory-mapped DVFS control register that allows the core to tune the DVFS mode and clock settings. The core and uncore logic is synthesized from an SDLS standard cell library using a standard automated place and route methodology, and uses static CMOS clock buffers to minimize insertion delay and improve slew rate relative to SDLS inverters. The list of the cells included in the SDLS library is shown in Fig. 3.12(a), and was kept small to minimize design time overhead. The cells were selected by first synthesizing the microprocessor core with a commercial standard cell library, and then choosing the cells that were mostly commonly used in the design. The core, uncore, and SRAM all use a 0.6 V nominal supply voltage.

The power and timing management subsystem operates at a 1.2 V supply and consists of a tunable reference current generator, VSC, and an ACG. The VSC generates V_{CN} and V_{CP} , each up to a maximum voltage of 1.0 V (for the nominal 0.6 V supply voltage) to overdrive V_{CN} and V_{CP} for a stronger cutoff effect. Fig. 3.12(b) shows the measured V_{CN} and V_{CP} outputs across V_{CSEL} , which is the memory-mapped register the core uses to control V_{CN} and V_{CP} . The ACG is also controlled by the core through a memory-mapped register, and is implemented using SDLS cells with the same V_{CN} and V_{CP} as the core, so that it can track the ripple and effectively replicate the critical path delay of the SDLS core.¹

3.4.4 Measured Results

The SDLS microprocessor test chip was fabricated in a 65 nm low-power CMOS process. This was chosen to demonstrate that DLS logic is relevant to newer technologies, as previous works

¹Further details on the implementation of the VSC and ACG are left to [JB10].

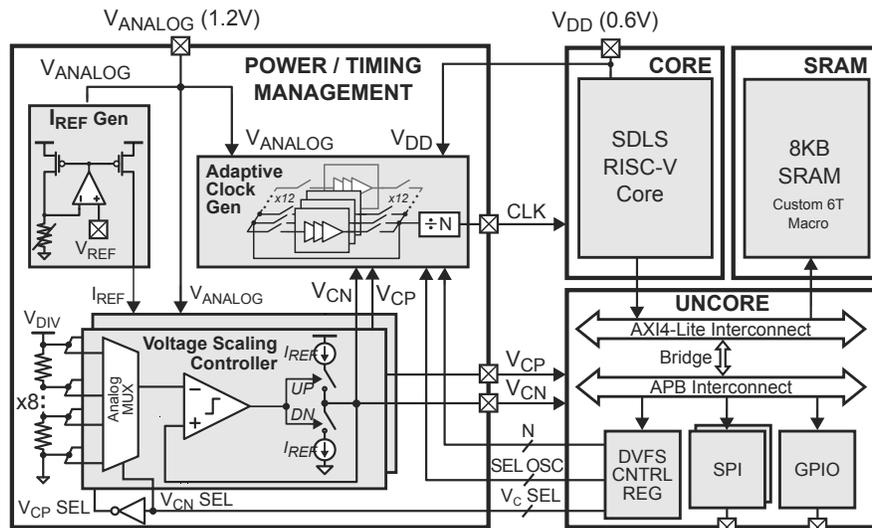
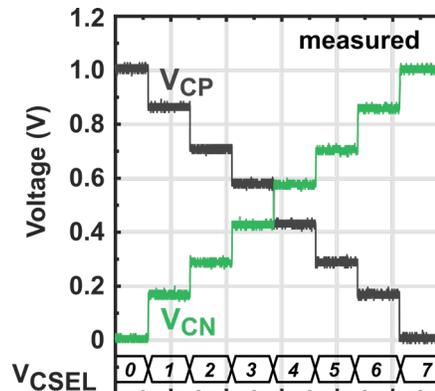


Fig. 3.11. System architecture of the proposed DVFS RISC-V microprocessor.

SDLS Library
AOI22
BUF
DFFRQ
INV
LATQ
MX2
NAND2
NAND3
NOR2
NOR3

(a)



(b)

Fig. 3.12. (a) The standard cells included in the SDLS library, and (b) the measured output of the VSC.

demonstrating DLS logic were implemented in 180 nm [35, 59] and 130 nm technologies [34]. Even newer technologies (i.e. FinFet), were not considered due to the associated cost and accessibility barriers. Fig. 3.13 shows oscilloscope measurements of V_{CN} , V_{CP} , the adaptive clock output, and the GPIO bits while the core executes a self-checking Fibonacci sequence program. While executing the program, the core dithers between two different V_{CSEL} modes, demonstrating that it can achieve an average power consumption that is between that of two different modes.

The program calculates all of the numbers in the Fibonacci sequence up to the maximum

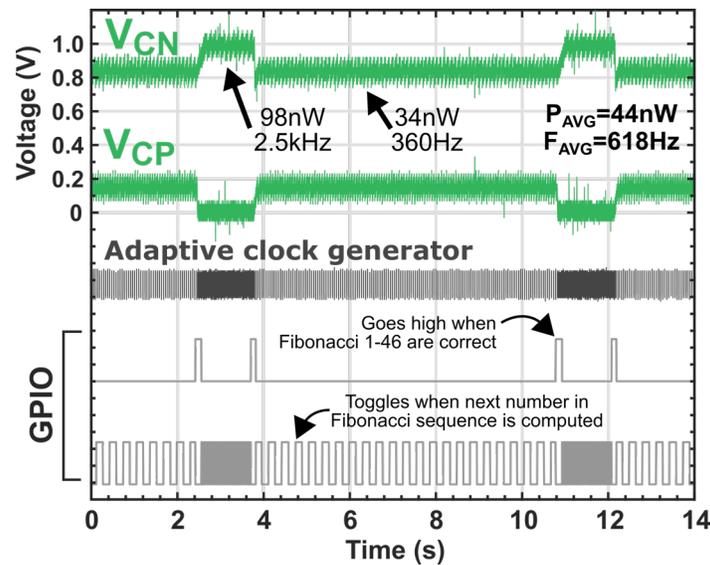


Fig. 3.13. Oscilloscope measurement of runtime DVFS, where the core dithers between two V_{CSEL} modes while computing the Fibonacci sequence. GPIO output bits indicate when a new Fibonacci number is computed and when the total sequence (numbers 1–46) are correct.

that can be represented with the 32 bit system architecture, and was designed to demonstrate that the core logic is fully functional. The last computed value in the sequence is compared with the correct value that is stored in the memory, and since each Fibonacci number is computed recursively from the previous number, this validates that all the previous calculations in the sequence were correct. The GPIO bits are used to indicate the progress of the program as it executes in a loop. The entire program, including the instructions that allow the core to control the voltage of V_{CN} , V_{CP} , and the adaptive clock frequency, were written in C and compiled with a standard open-source tool-chain.

To fully evaluate the SDLS core across supply voltage and DVFS mode, a simpler GPIO bit toggling program is used to more fairly compare to [34]. Fig. 3.14 shows the measured power of the core and uncore for each of the DVFS modes while running at the nominal V_{DD} of 0.6 V. Fig. 3.15 shows the measured maximum frequency and the minimum energy for each of the DVFS modes. At this 0.6 V V_{DD} , the SDLS core consumes 6 nW total power to run at 11 Hz in the minimum DVFS mode, and 140 nW total power at 8.2 kHz in the maximum DVFS mode. The uncore domain power increases from 4.86 nW to 93.9 nW under the same operating frequency conditions. Fig. 3.16 shows the full power breakdown at these minimum

and maximum DVFS modes, including measured power of the ACG and simulated power of the VSC, which consumes 3 nW total power from its 1.2 V supply across all V_{CSEL} modes.

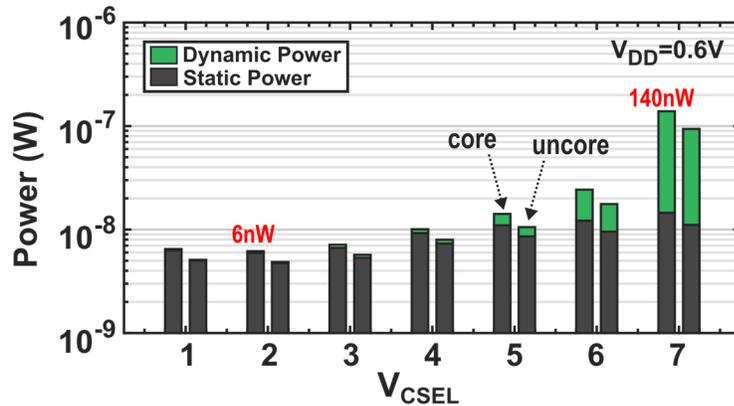


Fig. 3.14. The measured power of the core and uncore across DVFS modes at 0.6 V V_{DD} .

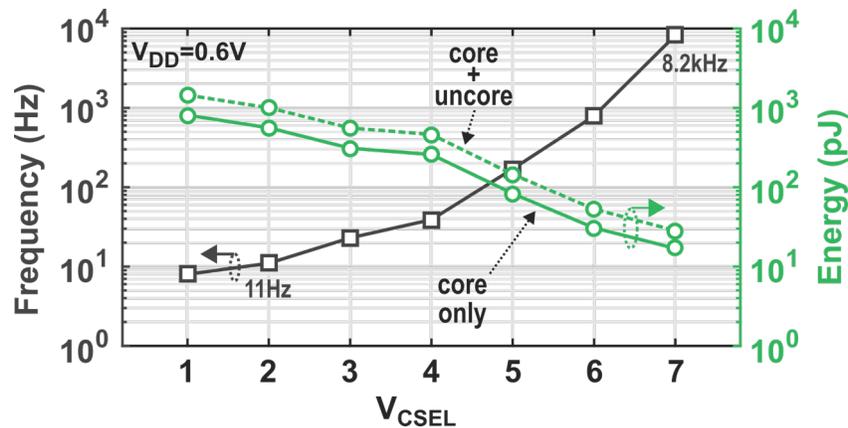


Fig. 3.15. The measured maximum frequency and energy across DVFS modes at 0.6 V V_{DD} .

Fig. 3.17 shows the achievable core power and frequency range from scaling the bias voltages V_{CN} and V_{CP} at each V_{DD} . The minimum achievable core power is 840 pW, which occurs in the minimum DVFS mode at 0.3 V V_{DD} while running at 6 Hz. At this point, the uncore domain consumes 690 pW and the clock generator consumes 130 pW, bringing the total minimum system power to 4.66 nW after the addition of the 3 nW VSC. While most processor cores require high operating frequency (and therefore, high power) to reduce energy consumption, the SDLS core achieves a competitive minimum energy of 13.4 pJ at 0.5 V V_{DD} while running at 2.07 kHz in the maximum DVFS mode, consuming just 27.9 nW. Increasing

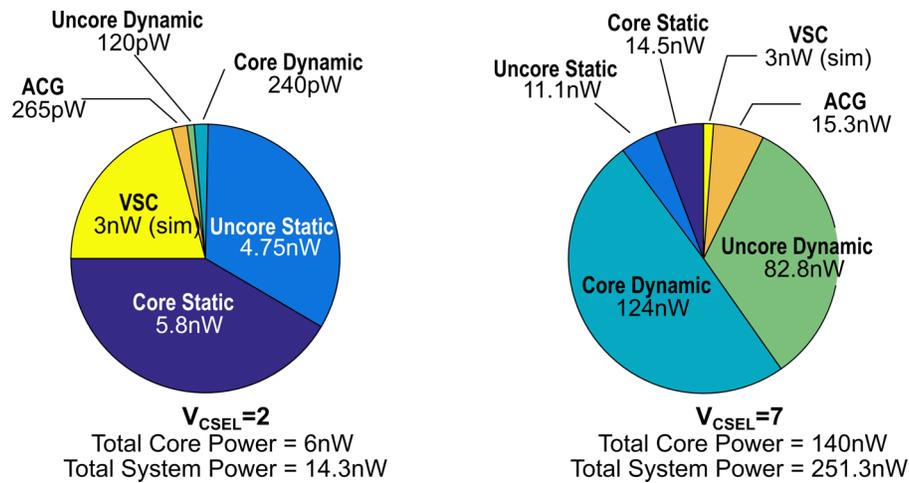


Fig. 3.16. The power breakdown for the minimum and maximum DVFS modes at 0.6 V V_{DD}

V_{DD} allows the core to reach higher frequencies of up to 41.5 kHz when $V_{DD} = 0.9$ V, $V_{CN} = 1.1$ V, and $V_{CP} = 0$ V. This results in high dynamic power though, due to the SDLS gates having a higher intrinsic gate capacitance than static CMOS gates.

3.4.5 Conclusion

This section presented a RISC-V microprocessor core in 65 nm CMOS, implemented with a SDLS logic family designed to enable constant V_{DD} performance scaling at ultra-low power levels. The included VSC and ACG enable a fully integrated modified DVFS scheme. At a constant 0.6 V supply voltage, the core can scale its power-performance from 6 nW at 11 Hz operating frequency to 140 nW at 8.2 kHz operating frequency. Across V_{DD} , the core achieves a minimum power of 840 pW, and a minimum energy of 13.4 pJ/cycle. Fig. 3.8 shows a comparison to state-of-the-art low-power processors, illustrating how this work enables a new area of the design space for continuous and reliable battery-less IoT sensing nodes. Fig. 3.19 shows a detailed summary of this work in comparison to other low-power microprocessor cores.

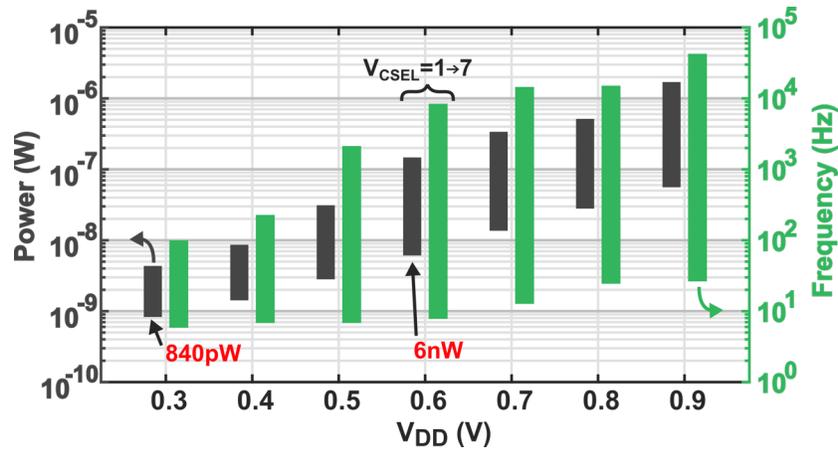


Fig. 3.17. Measured core DVFS tuning range versus V_{DD}

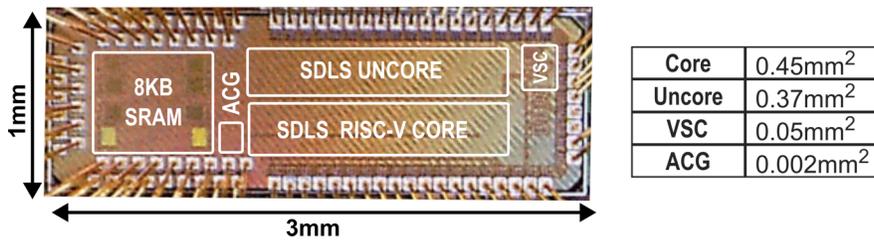


Fig. 3.18. Annotated die photograph of the SDLS RISC-V microprocessor.

TABLE I. PERFORMANCE SUMMARY AND STATE-OF-ART COMPARISON

	This Work		ISSCC '18 [8]	ISSCC '15 [7]	JSSC '17 [5]	ESSCIRC '18 [4]	JSSC '09 [2]
	Core	System**					
Technology	65nm		180nm	180nm	40nm	90nm	130nm
Architecture	RISC-V		MSP430	ARM Cortex M0+	ARM Cortex M0	ARM Cortex M3	8-bit
Area	0.45mm ²	0.87mm ²	5.33mm ²	2.04mm ²	0.16mm ²	7.18mm ²	0.83 mm ²
Operating Voltage	0.3V – 0.9V		0.2V – 1.1V	0.16V – 1.15V	0.2V – 0.5V	0.5V – 1.0V	0.45V – 0.9V
Performance Scaling	Modified DVFS		Dual-Mode	None	None	DVFS + AVS	None
Scaling Granularity	7-point + dithering		2-point	None	None	~40-point (12mV step)	None
Operating Frequency	6Hz – 41.5kHz		2Hz – 2.8MHz	2Hz – 15Hz	800kHz – 50MHz	1MHz – 20MHz	40kHz – 3MHz
Minimum Active Power	840pW @ 0.3V, 6Hz	4.66nW @ 0.3V, 6Hz	595pW* @ 0.45V, 2Hz	127.1pW* @ 0.55V, 2Hz	16.8μW @ 0.2V, 800kHz	34nW* @ 1.0V, 0Hz [†]	100nW* @ 0.45V, 40kHz
Minimum Energy	13.4pJ/cycle @ 0.5V, 2.07kHz	38.8pJ/cycle @ 0.5V, 2.07kHz	14pJ/cycle* @ 0.45V, 19kHz	44.7pJ/cycle* @ 0.55V, 7Hz	8.8pJ/cycle @ 0.37V, 13.7MHz	23pJ/cycle* @ 3V, 5MHz	2.8pJ/cycle* @ 0.35V, 106kHz

*includes memory, **System=core+uncore+ACG+VSC, †retention mode w/o power gating, ‡extracted from power breakdown

Fig. 3.19. Measured performance with comparison with the state-of-the-art low-power processors.

Chapter 4

Multi-Threshold DLS Logic

4.1 Motivation

Many IoT sensing applications have low sampling frequency requirements in the range of mHz to 100s of Hz, as previously shown in Table 1.2. The sampling frequency sets the lower bound on the clock frequency for the circuits performing these applications. Unless extensive on-chip processing of the sensor data is required, the clock frequency does not need to be significantly greater than the sampling frequency. Thus low-frequency circuits have many practical uses. In the sub-kHz region, dynamic power consumption is insignificant in static CMOS digital circuits, and static leakage power dominates. Section 3.1 discusses some common static power reduction techniques for static CMOS digital circuits, including standby modes, the use of H_{VT} devices, clock gating, and cutoff structures. As previously discussed in section 3.2, there is a need for digital circuits that can operate in an always-on mode at power levels below the leakage floor of conventional static CMOS circuits. Lower power consumption translates to more reliable operation from smaller energy harvesters, and enables self-powered IoT devices to operate in more energy-scare environments.

4.2 Prior Art

SDLS logic presented in section 3.4 enables always-on operation at power levels below the static CMOS leakage floor, but there isn't always a need to dynamically trade-off power and performance. SDLS logic also has significant overheads: the control voltages V_{CN} and V_{CP}

The content in this chapter is based on [JB13]. For a list of individual contributions see section 6.3.

(GIDL), and (2) $V_{DS,HN}$ being decreased, which reduces leakage due to Drain-Induced Barrier Lowering (DIBL). These two effects counteract, and result in a small net increase in I_{OFF} . A similar analysis can be done for the pull-down logic.

The L_{VT} FBB variant enables circuits to operate at up to a few kHz, but the leakage power is less than an order of magnitude below the leakage floor of static CMOS implemented with H_{VT} devices. Another DLS logic variant implemented in 65 nm using all IO devices with FBB was presented in [JB15]. IO devices utilize a thicker gate oxide than regular devices, so they have significantly less gate leakage, which limits the leakage power floor of non IO DLS logic in newer technologies such as 65 nm. This leads to the IO FBB variant being almost three orders of magnitude lower leakage power than static CMOS, but the maximum clock frequency is limited to a few Hz like the original DLS logic implementation in 180 nm CMOS [34, 59].

4.3 A Multi-Threshold Technique For Up To 6.5x Power Reduction in Dynamic Leakage Suppression Logic

4.3.1 Contributions

This section presents a gate-level Multi-Threshold technique for DLS logic in 65 nm CMOS that enables up to 6.5x power reduction compared to single-threshold DLS logic. Multi-Threshold Dynamic Leakage Suppression (MT-DLS) logic utilizes both the L_{VT} FBB and IO FBB DLS logic gates within the same circuit to bridge the power-performance gap between the two variants. Fig. 4.2 shows the 34x power gap between IO FBB and L_{VT} FBB DLS logic, the target application space for MT-DLS, and a comparison to H_{VT} static CMOS. MT-DLS logic yields power savings in the approximate 10 Hz–200 Hz range for a 0.4 V supply voltage, and in the 10 Hz–1.4 kHz range across the full 0.2 V–1.0 V supply voltage range.

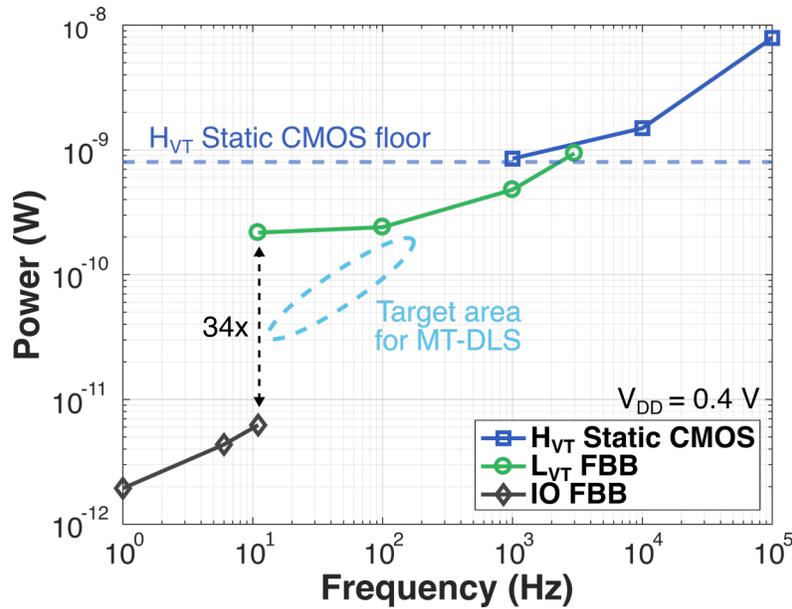


Fig. 4.2. Simulated power consumption of an 8-bit MAC block implemented with IO FBB DLS, L_{VT} FBB DLS, and H_{VT} static CMOS at 0.4 V supply voltage.

4.3.2 MT-DLS Logic Implementation

An 8-bit Multiply-Accumulate (MAC) block was chosen to evaluate MT-DLS logic since it is a simple sequential block used in many Digital Signal Processing (DSP) applications such as convolution, digital filtering, and Fast Fourier Transform (FFT). Fig. 4.3 shows the block diagram of the MAC block.

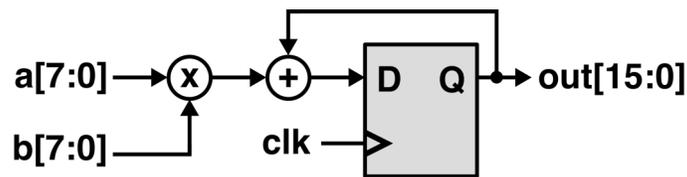


Fig. 4.3. Block diagram of the 8-bit Multiply-Accumulate (MAC) block.

A library of standard cells was created for both the L_{VT} FBB and IO FBB DLS logic variants, and Fig. 4.4(a) shows the cell list. For fair comparison, an H_{VT} static CMOS library limited to the same cells was also created from a larger commercial library. Timing libraries were created for each of the DLS variants, and synthesis scripts were developed to enable their simultaneous use within the same circuit. The synthesis scripts were designed to prefer IO

FBB cells to minimize the power consumption, and use only the minimum number of L_{VT} FBB cells needed to satisfy the clock frequency constraint. Special tap-cells were also created to body bias the DLS cells, and a standard Automated Place and Route (APR) flow was used to implement the MAC blocks. Both DLS variants share the same cell footprints, thus MT-DLS doesn't significantly impact circuit area. Fig. 4.4(b) shows an area comparison between the final placed-and-routed DLS and static CMOS MAC blocks.

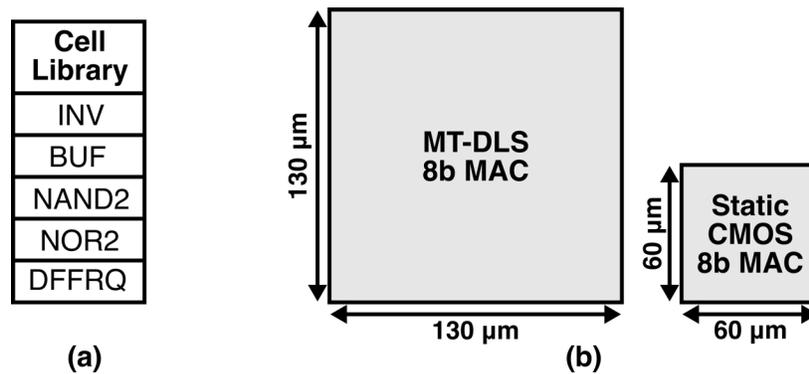


Fig. 4.4. (a) The DLS standard cell library used in this section. (b) Area comparison between DLS and static CMOS for the 8-bit MAC block.

4.3.3 Simulated Results

An array of 5 MAC blocks were created with MT-DLS logic, and designed to target different trade-offs between power and performance. The MAC blocks were tested with an input vector of 100 randomly generated inputs, and if all 100 outputs were correct, the block was considered functional for the operating conditions (V_{DD} , frequency, etc.). All simulations were performed at 27 °C and Typical Typical (TT) process corner unless otherwise noted. Fig. 4.5 shows that at 0.4 V supply voltage, MT-DLS provides power savings of up to 6.5x, and power savings throughout the approximate 10 to 200 Hz region. The listed cell percentages are derived from the gate count, with all logic gates normalized to the area of the 2-input NAND gate. Fig. 4.6 shows the power consumption and max frequency of the MT-DLS MAC block implemented with 65% L_{VT} FBB cells across the supply voltage range. DLS logic in 65 nm exhibits a power minimum at approximately 0.4 V due to competing trends in Sub-Threshold (Sub-Vt) leakage

and gate leakage. Sub-V_t leakage decreases with V_{DD} due to the more negative V_{GS,HN} (or V_{SG,FP}), but gate leakage increases due to its quadratic relationship with V_{DG,HN} (or V_{GD,FP}), which is equal to V_{DD} in the steady state.

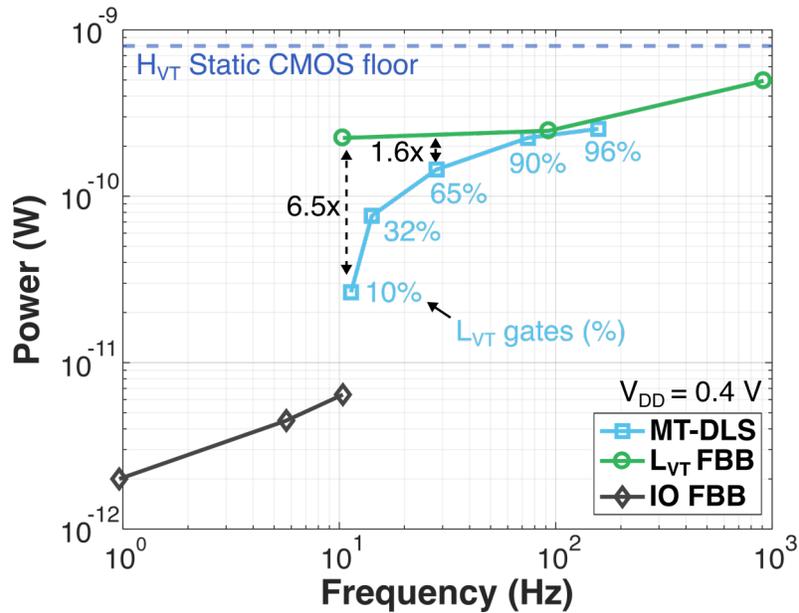


Fig. 4.5. Simulated power consumption of the MT-DLS 8-bit MAC blocks at 0.4 V supply voltage.

Fig. 4.7(a) shows the relationship between leakage power and the percentage of L_{VT} gates used in the MAC blocks, which is approximately linear since the IO FBB gates are 2 orders of magnitude lower leakage than the L_{VT} FBB gates. Fig. 4.7(b) shows the relationship between the max frequency and the percentage of L_{VT} gates used in the MAC blocks. The shape of this curve is determined by the timing path length distribution, and is thus highly circuit specific. For the MAC block in this work, there is one set of critical timings paths from the inputs a/b to the output register, so the timing path length distribution is relatively narrow. Larger more complex circuits typically have a wider timing path length distribution, with more non-timing critical paths where the slower logic gates can be used without effecting the max clock frequency. Therefore, larger circuits will likely benefit significantly more from MT-DLS than the simple MAC block presented in this work.

Fig. 4.8 shows the power reduction enabled by MT-DLS across the clock frequency and supply voltage range. MT-DLS yields power savings of up to 85% at 0.4 V supply voltage

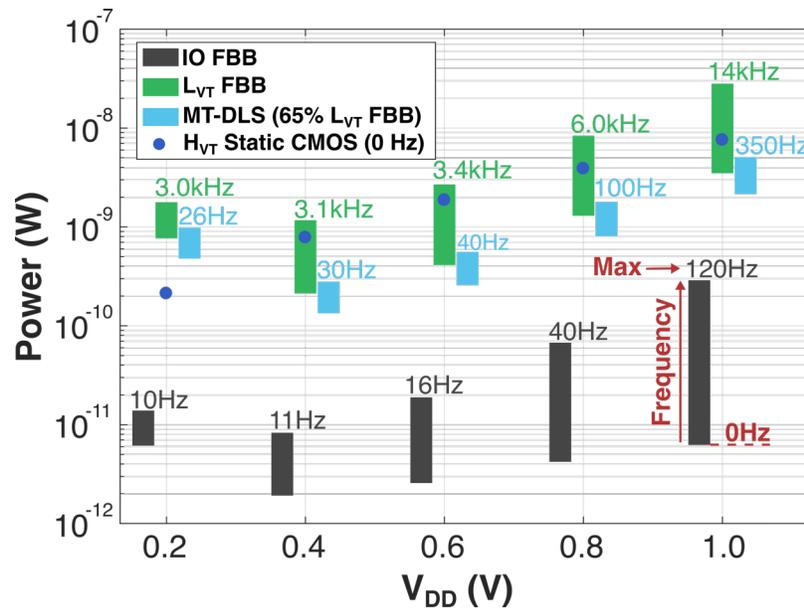


Fig. 4.6. Simulated power consumption and max clock frequency across the 0.2 V–1.0 V supply voltage range for the 8-bit MAC block implemented using L_{VT} FBB DLS logic, IO FBB DLS logic, and MT-DLS logic with 65% L_{VT} FBB logic gates.

compared to L_{VT} FBB logic, and power savings throughout the 10 Hz–1.4 kHz range across the 0.2 V–1.0 V supply voltage range.

4.3.4 Conclusion

This section presented a gate-level Multi-Threshold technique for DLS logic (MT-DLS) in 65 nm CMOS, which combines L_{VT} FBB and IO FBB DLS logic gates within the same circuit to bridge the power-performance gap between the two variants, and enables up to 6.5x power savings. MT-DLS yields power savings throughout the 10 Hz–1.4 kHz range across the 0.2 V–1.0 V supply voltage range. Ultimately, this improves the state-of-the-art power-performance trade-off for low-frequency digital circuits that operate below the leakage floor of conventional static CMOS, and helps enable self-powered IoT devices to operate in more energy-scare environments.

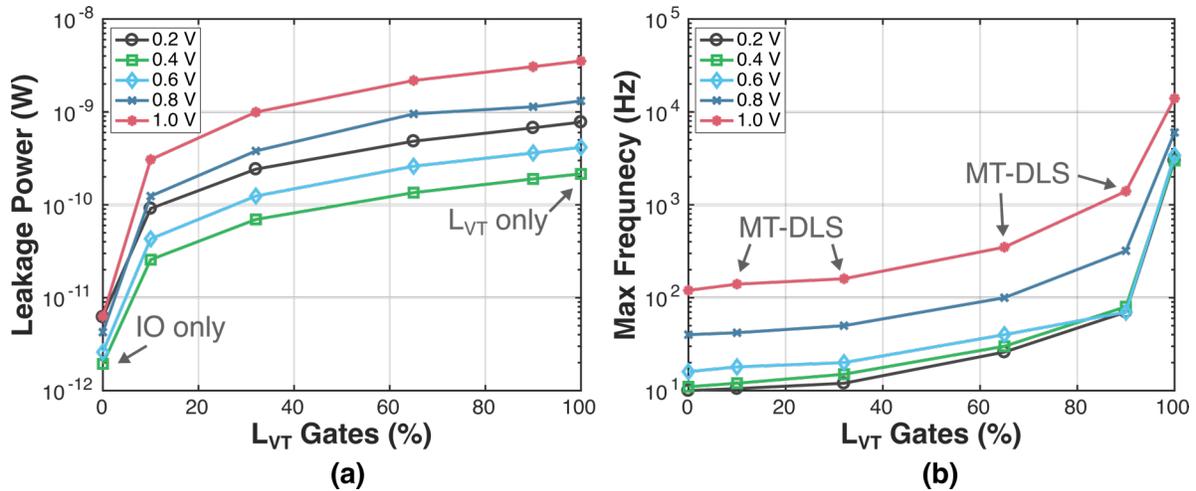


Fig. 4.7. Simulated leakage power consumption and maximum clock frequency for the 8-bit MAC blocks implemented with varying percentages of L_{VT} FBB and IO FBB DLS logic gates.

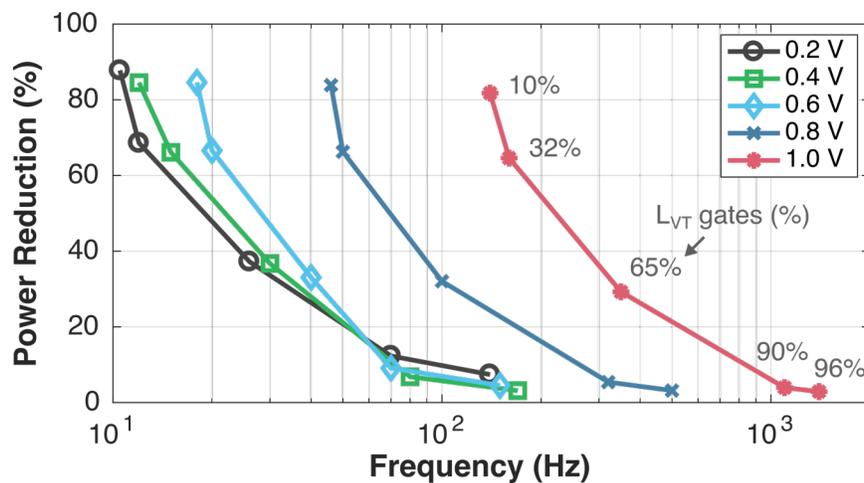


Fig. 4.8. Simulated power reduction for the 8-bit MAC block enabled by MT-DLS across clock frequency and supply voltage. Power reduction is determined by comparison to the MAC block implemented with all L_{VT} FBB logic.

Chapter 5

Ultra-Low Power IoT Systems and Applications

5.1 Motivation

Commercial products that utilize ULP and self-powered IoT devices exist [30–32], but they're not prolific or widely adopted yet. This is partly because the value of self-powered operation is not widely understood yet, as it only manifests in very large-scale deployments. The overhead of battery recharge and replacement for a deployment with only a few thousand devices isn't significant if batteries last months to years. Another reason is that it's technically difficult to achieve reliable self-powered operation with a small device size. Chapters 2-4 of this dissertation attempt to address this by lowering the power consumption of wireless TXs and digital processing circuits. But the latest advancements in technology must still be translated into functional system prototypes, and applications with commercial value must be demonstrated before widespread adoption will occur.

This chapter presents two systems and their applications for the purpose of accelerating the commercial adoption and proliferation of IoT devices. First, we present a self-powered SoC [JB3,JB7] and use it to implement several health-monitoring applications, including free fall detection, temperature sensing, and arrhythmia detection using the on-chip AFE [JB4]. Then, we integrate the SoC into a vigilant cardiac monitoring system that is powered completely by body heat, and tightly integrate the system into a compact wearable form to enable comfortable long-term use [JB14]. We also present an ULP wireless wake-up and control system, which is designed to reduce the electricity usage of MELs, and includes a fully integrated 802.11ba

The content in this chapter is based on [JB3, JB4, JB7, JB11, JB12, JB14]. For a list of individual contributions see section 6.3.

WuRX [JB12] and an 85 nW node-controlling SoC [JB11].

5.2 IoT SoC for Health-Monitoring Applications

The applications presented in this section are implemented with the health-monitoring SoC presented in [JB3,JB7]. Fig. 5.1 shows the system architecture. The SoC is implemented in a 130-nm CMOS technology, and features an array of sensing interfaces and DSP accelerators targeting health-monitoring and general-purpose IoT applications. It includes a custom 16-bit microprocessor core (LPC), two integrated 2-kB SRAMs for instruction and data memory, and a crystal-based 32 kHz oscillator for the digital subsystem. There is also an on-chip energy harvester that can harvest from either a TEG or PV cell, and a Power Management Unit (PMU) that generates regulated supply voltages for the SoC from an unregulated energy storage device.

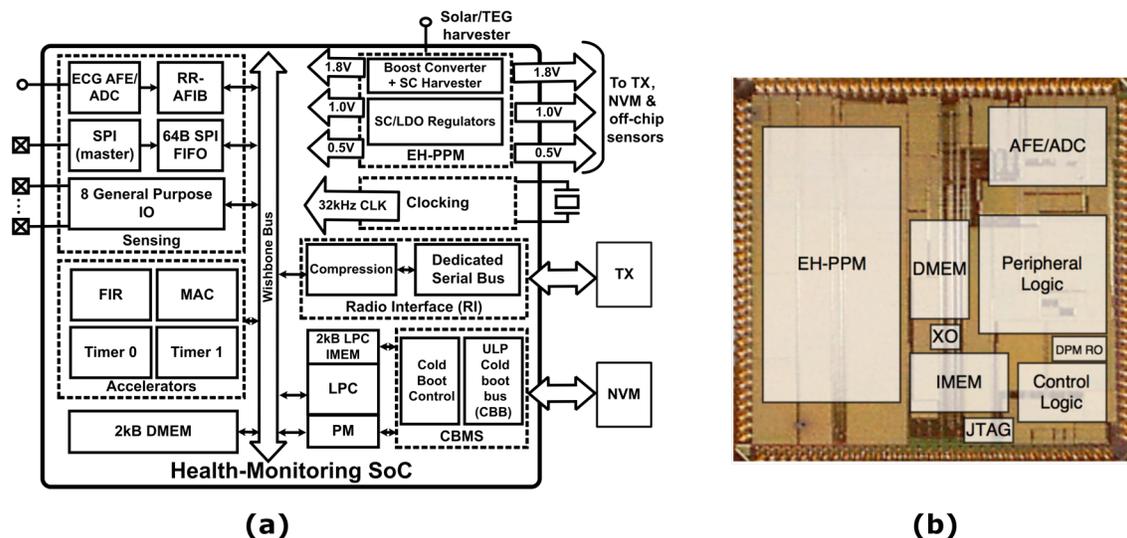


Fig. 5.1. (a) Block diagram of the health-monitoring SoC, and (b) the annotated die photograph.

5.2.1 Free Fall Detection

At rest, objects on the surface of the earth experience a constant acceleration of 9.81 m s^{-1} , or 1 g. A free fall occurs when an object falls through space without significant resistance, and experiences approximately 0 g of acceleration. Detecting free falls has a number of useful

applications. It can be used for commercial applications like shipping and handling impact tracking, where the magnitude and number of falls a package experiences is monitored over time. If the magnitude or number of falls exceeds some threshold, the customer can then avoid accepting a damaged shipment that was improperly handled. Such commercial products already exist [62], but electronic versions are bulky and not self-powered. And non-electronic versions do not provide detailed information about the exact magnitude or number of falls. Free fall detection can also be used for detecting falls in humans, a health-monitoring application. Falls are the leading cause of fatal and nonfatal injuries among persons aged 65 or greater, and in 2014, 28.7% of older adults reported falling at least once in the preceding 12 months [63]. Often, the most serious consequences of a fall are due to a delay in treatment, as the injuries sustained during the fall prevent the person from obtaining help. An IoT device could be used to detect falls and send a signal for help, but it must be self-powered to eliminate the user burden of battery recharge and replacement. It must also be small enough that it can easily be integrated into clothing or wearable accessories without significant discomfort, which would reduce user acceptance.

Micro-Electro-Mechanical System (MEMS) based accelerometers are commercially available with ultra-low power consumption. The ADXL362 accelerometer from Analog Devices consumes just $3.6\ \mu\text{W}$ at 2.0 V when sampling acceleration data at 100 Hz [64]. It is just a MEMS sensor with an Analog to Digital Converter (ADC), and a simple digital interface though. It cannot perform a complete application like free fall detection on its own.

In this section, we integrate the ADXL362 accelerometer with the health-monitoring SoC and a wireless TX to form a complete and fully self-powered sensor system that can perform free fall detection. The application flow diagram is shown in Fig. 5.2. First, the SoC configures the accelerometer for free fall detection using the SPI. The accelerometer is programmed to assert one of its interrupt pins when the acceleration falls below 0.6 g for a 200 ms period of time. We chose this simple threshold-based approach to minimize the system power. Although more sophisticated methods are available [65] and can give better results by preventing false

positives, they require more processing of the acceleration data. The interrupt pin of the accelerometer is connected to one of the SoC GPIO pins, and wakes the SoC up from a stalled state, in which parts of the memory are switched off to reduce power. When the interrupt is triggered, the SoC switches on the SPI block and reads the acceleration data surrounding the free fall that was logged by the accelerometer in its First-In First-Out (FIFO) buffer. The data is then compressed using the on-chip accelerator [JB2] to minimize the amount of data that must be transmitted, and thus reduce the wireless TX power. The TX is then switched on, and the data is sent before the TX is switched off again, and the system resets to detect the next free fall occurrence.

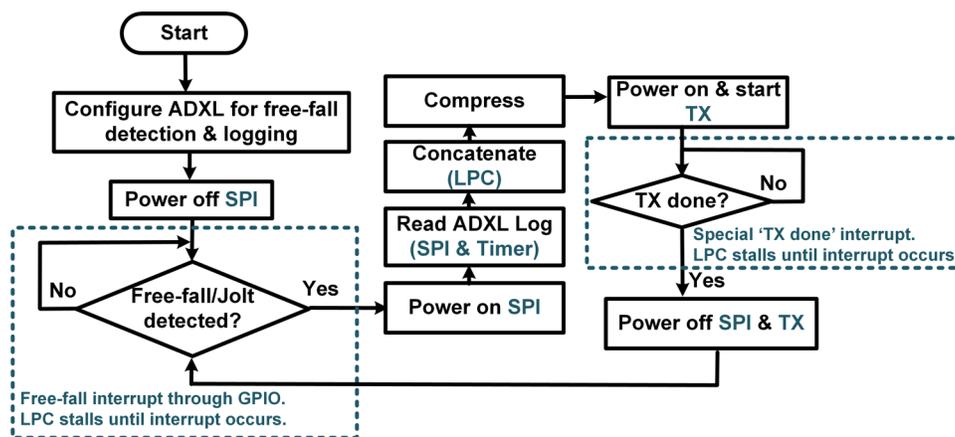


Fig. 5.2. Flow diagram for the proposed free fall detection application.

Fig. 5.3(a) shows the acceleration data that was captured by the SoC for a free fall event. Fig. 5.3(b) shows the SoC power breakdown for this application. The active state is defined as the part of the application where the SoC is reading the data from the accelerometer, compressing it, and transmitting it using the off-chip TX. Compression reduces the active SoC power without even accounting for the TX power reduction, as less data must be sent through the pads to the off-chip TX. With a pessimistic assumption of 1 free fall per minute, the SoC power consumption is 507 nW. The entire system consumes just 20.6 μ W when measured from the energy storage super-capacitor, including the voltage regulation overhead, accelerometer, and wireless TX.

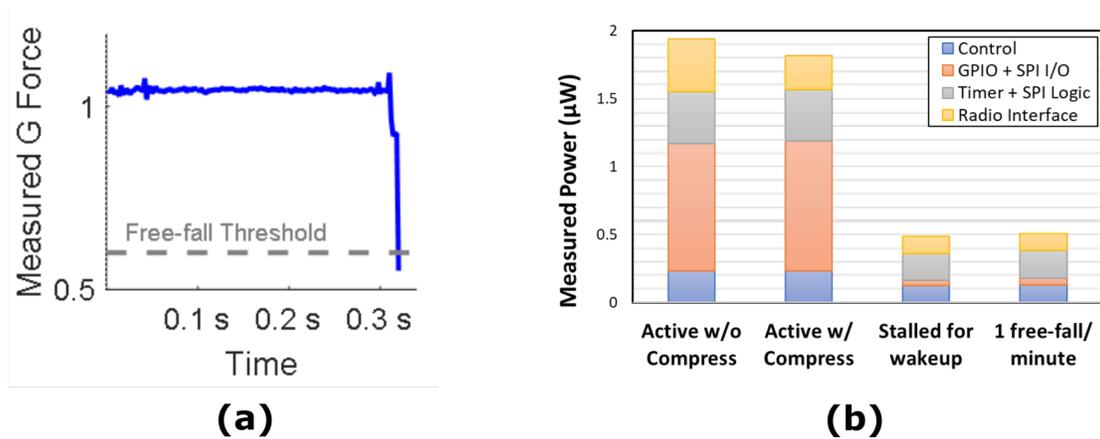


Fig. 5.3. (a) Acceleration data captured by the SoC for a free fall. (b) Measured SoC power consumption for the free fall detection application.

5.2.2 Temperature Sensing

The ADXL362 accelerometer also includes a temperature sensor for calibrating the MEMS sensor as environmental conditions change. To further characterize the SoC's performance for IoT applications with varying sampling rate requirements, we created a simple application where the SoC reads and transmits the ambient temperature from the accelerometer once every 5 minutes. For this application, the SoC consumes $2.06 \mu\text{W}$ when actively reading and transmitting the temperature data, $0.518 \mu\text{W}$ when waiting between samples, and $0.519 \mu\text{W}$ overall. This illustrates that for low sampling frequency applications, the overall SoC power is set by the standby mode power. The total system consumes $7.61 \mu\text{W}$ on average when measured from the energy storage super-capacitor, including the accelerometer and TX.

5.2.3 Arrhythmia Detection

Global health care expenditures are increasing at an alarming rate [66], and cardiovascular diseases are the leading cause of death globally, accounting for 31% of all deaths in 2016 [67]. Many cardiac conditions are difficult to detect in short periods of time, and thus require long-term continuous monitoring. Although commercial wearable cardiac monitoring systems exist [68], they are typically only used for short periods of time on the order of days due to their size and imposed user burden. Making these systems self-powered reduces the user burden of

battery recharge and replacement, and thus improves user compliance. And reducing the power consumption further reduces the required size of the energy harvester and storage device, making the device smaller and easier to wear for long periods.

This section presents a ECG acquisition and arrhythmia detection application, implemented with the health-monitoring SoC, its integrated AFE [JB4], and 12-bit Successive Approximation Register (SAR) ADC. The SoC first configures the ADC to sample data at 100 Hz and to generate an interrupt when a new sample is available. In the interrupt service routine, the microprocessor core moves the data from the ADC to the radio interface, where the data is compressed and buffered until enough is available to fill a packet, and then finally transmitted. Fig. 2.6(a) shows the ECG signal captured by the SoC. The average power of the SoC is $1.02 \mu\text{W}$, and the power breakdown is shown in Fig. 5.4(b). The total system power measured from the energy storage node is $5.98 \mu\text{W}$, including the voltage regulation overhead and wireless TX.

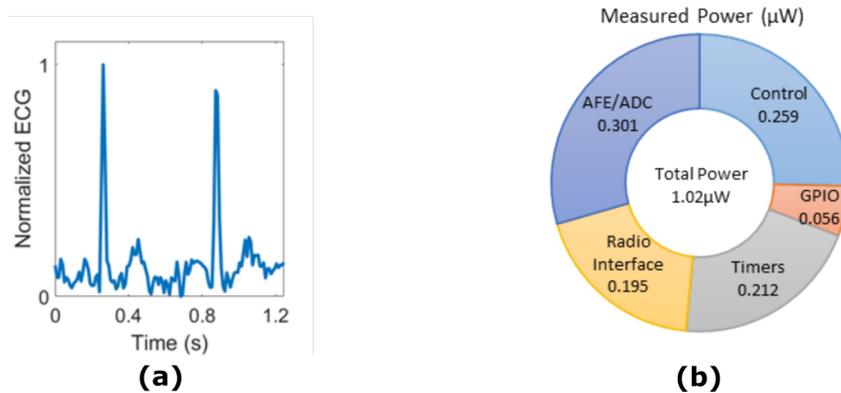


Fig. 5.4. (a) ECG data captured by the SoC AFE and ADC. (b) Measured SoC power consumption for the ECG capture application.

To validate that the acquired ECG signal is clinically useful, the data was processed off-chip using a heart rate extraction and arrhythmia detection algorithm.¹ The R-R (heart rate) extraction algorithm is based on the popular Pan-Tompkins algorithm [69]. The algorithm uses an initial 4 s time window to estimate the DC baseline, and then uses amplitude thresholds and

¹The on-chip block designed to perform this application was non-functional due to a design bug, so the same algorithms were implemented off-chip to validate the data. I was not involved in the design or implementation of the on-chip block.

time windows to extract the R peaks. Fig. 5.5 shows the measured ECG signal, sampled at 200 Hz, with annotations indicating the number of samples between consecutive R peaks. The arrhythmia detection algorithm uses the number of samples between consecutive R peaks, and is able to detect AFib in as little as 12 heart beats [70]. Two irregular missed heart beats shown in Fig. 5.5, and are detected by the algorithm as an AFib event.

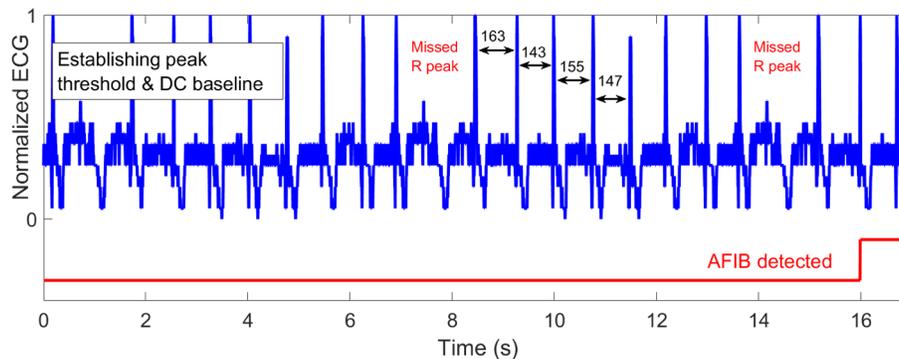


Fig. 5.5. The captured ECG signal, showing AFib being detected. The annotations indicate the number of samples between consecutive R peaks.

5.2.4 Wearable Cardiac Monitoring System Powered by Body Heat

This section presents a vigilant cardiac monitoring system that is completely powered by body heat, and tightly integrated into a compact wearable form to enable comfortable long-term use. In this context, vigilant means the system operates in a mode that is always active, and no critical AFib events are missed due to intermittent operation. Fig. 5.6 shows an overview of the system, which integrates numerous state-of-the-art components developed by multiple universities. The system spans three Printed Circuit Boards (PCBs): a primary board with the health-monitoring SoC, a radio board with a Bluetooth Low Energy (BLE) compliant TX, and an energy harvesting board with a super-capacitor and DC-DC converter. The boards were all designed for minimum area, and are integrated vertically into a 3D-printed case.

The case attaches magnetically to the back of a custom e-textile shirt, shown in Fig. 5.7. The shirt integrates dry ECG electrodes, which operate without gels or adhesives that can irritate the skin, and thus greatly improve user comfort during long-term use. The shirt is constructed

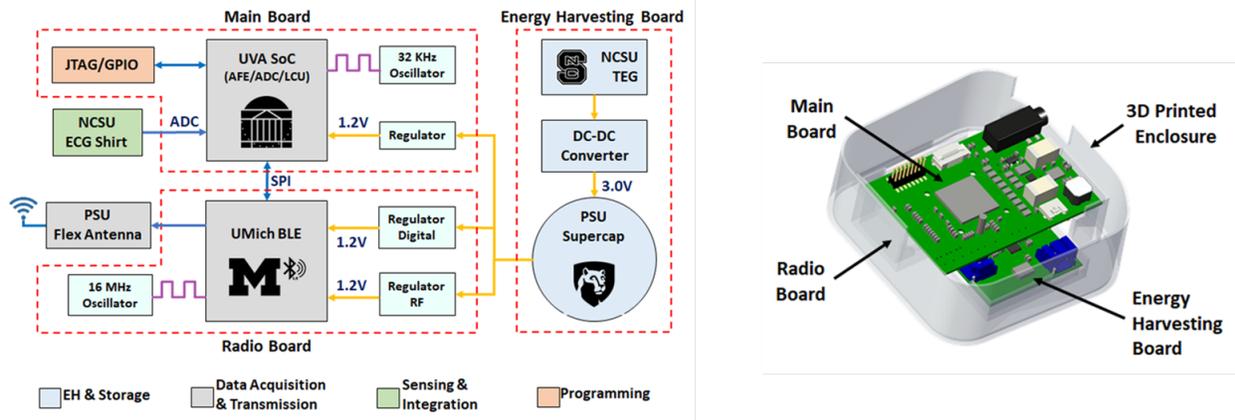


Fig. 5.6. Block diagram of the proposed three PCB system, and the PCBs integration into the 3D-printed case.

with multiple fabrics, and uses high compression panels to ensure that the electrodes make good connection with the skin. The shirt also integrates a panel of flexible TEGs on the front, which powers the system using the temperature differential between the user's skin and the air. The TEGs connect to the DC-DC converter on the energy harvester board, and any excess energy is stored on an ultra-low leakage 0.92 F super-capacitor. A flexible antenna is also integrated into the shoulder of the shirt for the BLE TX, and is designed specifically to cope with the effects of human body loading presented in wearable applications.

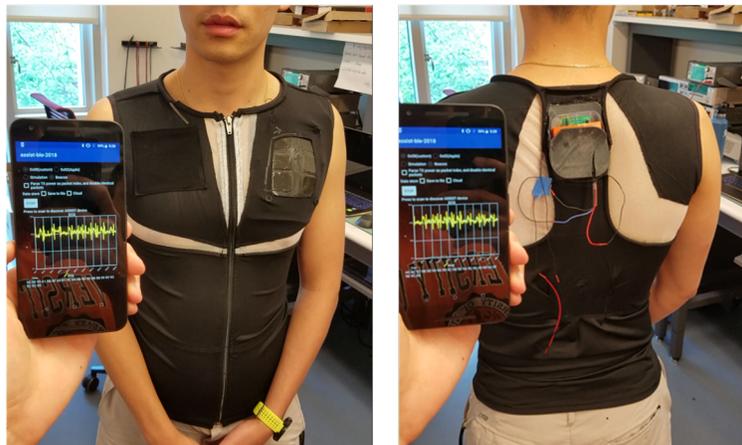


Fig. 5.7. The proposed wearable system is integrated into a custom e-textile shirt, and streams ECG data in real-time using BLE to a smartphone, from which the data is then uploaded to a cloud system.

The SoC samples ECG data at 50 Hz and only uses the 8 most significant bits of the 12-bit ADC to save power, as this was previously shown to be the minimum needed for vigilant

operation as defined by AFib detection [71]. The TX transmits the ECG data in real-time to a smartphone, where it is displayed on a custom Android app, shown in Fig. 5.7. From the smartphone, the data is uploaded to a cloud system where it can be remotely viewed by healthcare providers. The average power of the entire system is $65 \mu\text{W}$ at the power neutral point, which is defined as the super-capacitor voltage where the harvested power from the TEGs is equal to the power consumed by the system from the super-capacitor. This was measured with the user standing still indoors, and not moving to generate additional airflow across the TEGs. Fig. 5.8 shows the system power consumption as a function of the super-capacitor voltage. The super-capacitor voltage is 2.15 V at the power neutral point, and when the super-capacitor is fully charged to 3 V , the system can operate for over 9 hours without any harvesting, as shown in Fig. 5.8(b).

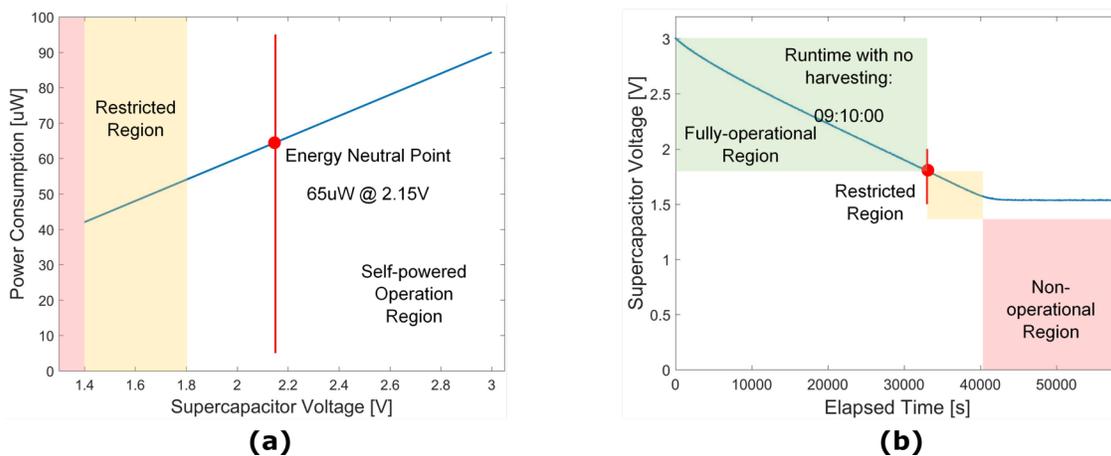


Fig. 5.8. (a) The system power consumption as a function of the super-capacitor voltage, and (b) the super capacitor voltage as a function of time with no harvesting.

5.3 IoT SoC for MELs Power Management

5.3.1 Motivation

MELs in buildings are electric loads resulting from devices such as power adapters, televisions, and coffee makers [72]. MELs do not include major appliances such as refrigerators, stoves, dryers, etc., or other major electronic systems such as those responsible for heating, cooling,

and lighting. Nonetheless, in 2016 MELs accounted for 30% of electricity usage in residential buildings, and 36% in commercial buildings according to the U.S. Department of Energy (DoE) [73]. A significant part of the problem is that many MELs consume power even when they are switched off, often on the order of magnitude of watts [74]. This is often referred to as phantom standby energy consumption, and it is becoming an increasingly significant problem for the future. The number of MELs is rapidly growing, and many devices that have traditionally consumed zero energy in their off state are being replaced with new smart devices, which have nonzero standby energy consumption. This energy isn't entirely wasted though, as it typically provides some useful function. One example is the Philips Hue smart Light-Emitting Diode (LED) light bulb, which uses energy in the standby state to allow itself to be controlled by smartphone via Bluetooth. According to the manufacturer, it has a max standby power of 0.5 W, and a max active power of 17 W [75]. Thus in some instances, these smart LED bulbs can result in greater total energy use than traditional incandescent or fluorescent bulbs. This occurs when the smart bulbs are not actively used much, and their total power is dominated by the standby power, which is zero for traditional non-smart light bulbs.

New strategies must be derived to address the growing power consumption of MELs and their nonzero standby energy consumption. Although MEL standby energy is typically used to perform a useful function, the always-connected always-ready standby state employed by many emerging smart devices isn't required at all times. For example, if a building is empty, or its occupants are not near the smart device, it could be switched off completely to eliminate its standby energy consumption without inconveniencing the occupants. To accomplish this, information collected from sensors distributed throughout the building and the mobile devices of its occupants must be aggregated, and new algorithms to process the data must be invented. Furthermore, new devices must be created to manage the power states of MELs, and optionally disconnect them from the grid to eliminate their standby energy use. These devices must have ultra-low power consumption, orders of magnitudes less than the standby power of typical MELs in order to effectively address MEL standby energy consumption. They must also have wireless

connectivity that is standards compliant, so that they can be easily deployed and integrated into existing networks.

Wi-Fi is ubiquitous and found in almost every building, but commercially available Wi-Fi receivers consume > 100 mW when active [76]. WuRXs are an emerging technology, and are recently becoming integrated into existing standards [77]. They consume orders of magnitude less power than traditional receivers, and can be used to duty-cycle and wake up on demand traditional high-power receivers. They are therefore an ideal choice for enabling wireless connectivity in an ULP MEL control device. In addition to wireless connectivity, a MEL control device must also be able to perform basic computations and execute simple algorithms to control the power states of MELs. Commercially available low-power microcontrollers such as the popular MSP430 from Texas Instruments consume > 1 mW active power though [78], and are therefore not a viable option.

5.3.2 Contributions

This section proposes an ULP wireless wake-up and control system, designed to manage the power states of MELs when they are being used, and also to disconnect them from the electric grid when they are not being used to eliminate their phantom standby energy consumption. Fig. 5.9 shows an overview of the system. It includes of a fully integrated 0.2 V 802.11ba WuRX [JB12] and an 85 nW active power node-controlling SoC [JB11], which is the focus of this section.

In this system, an external gateway device aggregates inputs from sensors distributed throughout the building and the occupant's mobile devices, and forwards the relevant information for MEL control to the system. The WuRX consumes 578 μ W in its active state, and receives 32-bit wake-up codes that are decoded by the 16-channel correlator on the SoC. The SoC's multiple channel correlator and 32-bit RISC-V microprocessor core allows the SoC to control multiple MELs at once, and execute algorithms to control MEL power modes. This is done using an off-chip relay and de-multiplexer device, which interfaces to the MELs, and can optionally

disconnect them from the grid to eliminate their phantom energy.

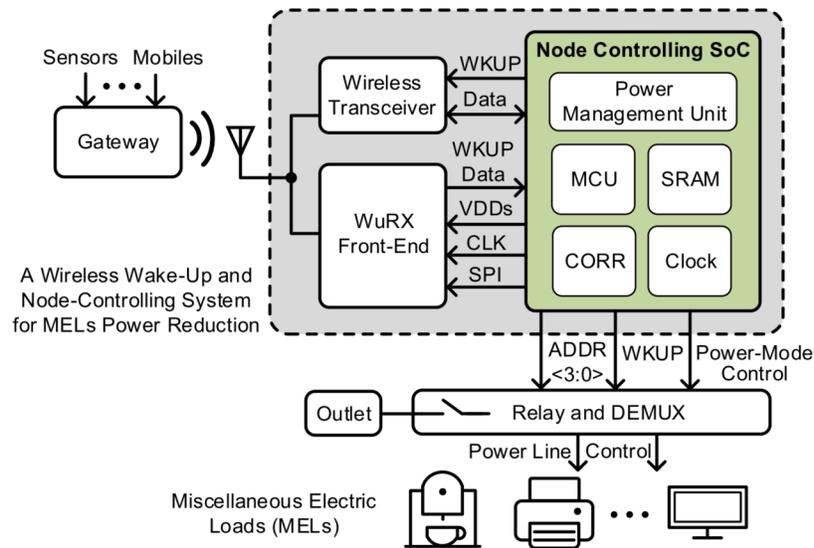


Fig. 5.9. The proposed wireless wake-up and control system for addressing MEL power consumption.

5.3.3 System Architecture

The SoC is implemented in a 65 nm low-power CMOS technology, and the system architecture and die photograph are shown in Fig. 5.10. The SoC includes a BottleRocket 32-bit RISC-V microcontroller-class processor core, an integrated 8-kB SRAM, and SPI and GPIO peripherals for interfacing to the WuRX and MELs. It also includes a correlator that samples the serial data output from the WuRX, and compares it to 16 different programmable 32-bit wake-up codes stored on the SoC to look for a match. The correlator is implemented in its own clock domain to reduce the system power, as it requires a higher frequency to sample the WuRX output than the rest of the digital subsystem, which operates at just 1 kHz. The digital subsystem was designed to operate at 0.45 V in the subthreshold region, and uses standard cells with H_{VT} devices to minimize the leakage and total power consumption. The SoC also includes a digital buck converter and several Low-Dropout (LDO) regulators, which generate regulated supply voltages for the digital subsystem and the WuRX. The entire wireless wakeup and control system can thus be powered with a single supply voltage, generated by a battery or an energy harvester.

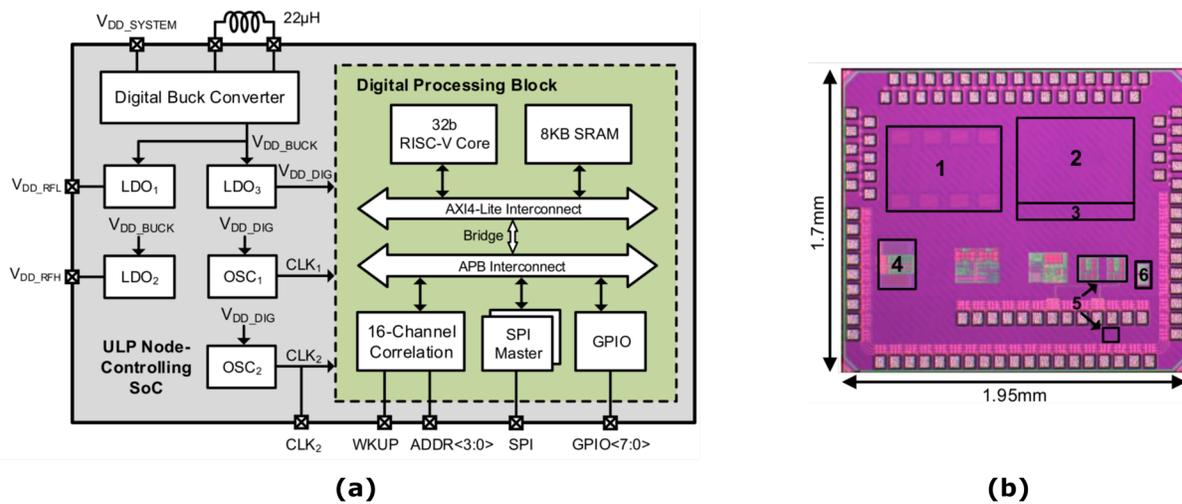


Fig. 5.10. (a) The proposed node-controlling SoC system block diagram, and (b) the SoC die photograph.

5.3.4 Measured Results

To demonstrate the SoC’s capability to control the power states of MELs at ultra-low power levels, we created an example application, shown in Fig. 5.11. In this application, the SoC uses GPIO to interface to an off-chip relay that controls the power states of MELs, and also an off-chip radio transceiver. Essentially, this is a GPIO bit-banging application that follows the state transition diagram shown in Fig. 5.11. The wake-up signal from the WuRX, software-based timers, and some GPIO input bits are used to transition the state, which is represented by two of the GPIO output bits. Fig. 5.12 shows the measured GPIO bits, with annotations indicating the state transitions.

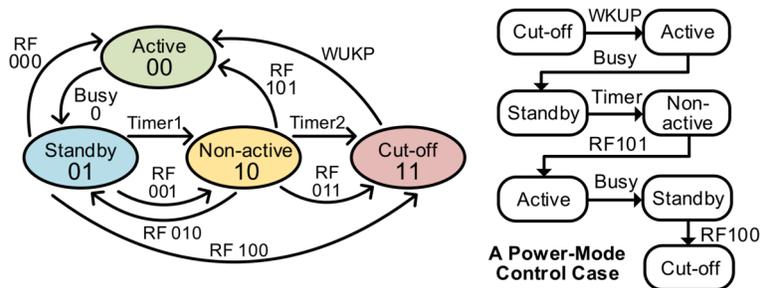


Fig. 5.11. State diagram for the example application, and flow chart showing the MEL operating mode transitions

For this application, the total SoC power is 85 nW from the 1 V input supply, with the digital subsystem operating at 0.45 V and 1 kHz. Fig. 5.13 shows the power consumption of the digital

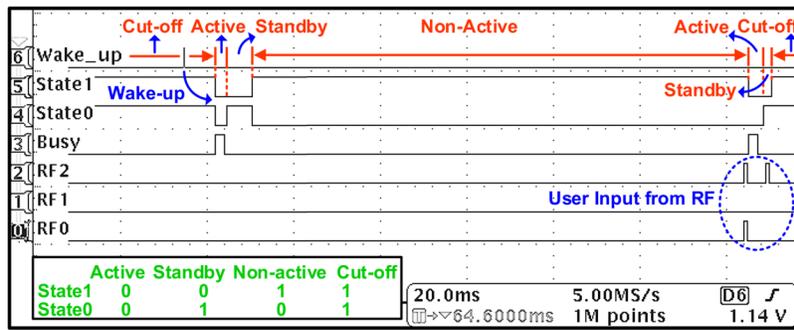


Fig. 5.12. Measured waveform showing wake-up signal, and the GPIO bits that represent the MEL power mode, MEL busy status, and the RF transceiver input

subsystem across supply voltage and frequency, excluding the correlator which is on a separate supply. We also tested the SoC with the WuRX, shown in Fig. 5.14. The SoC generates the regulated voltage supplies for the WuRX, controls it using the SPI, and is able to received a wake-up signal.

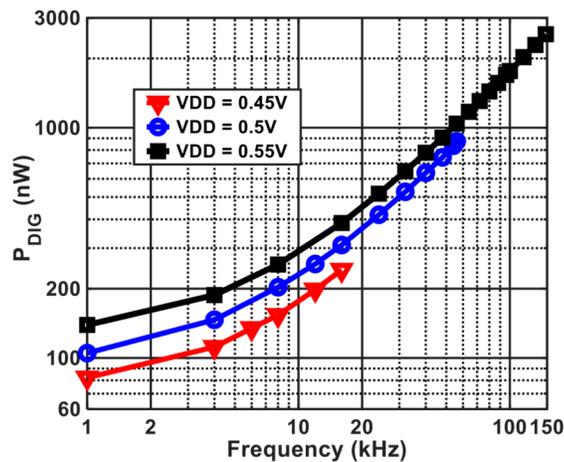


Fig. 5.13. The SoC digital subsystem power consumption as a function of operating frequency and supply voltage.

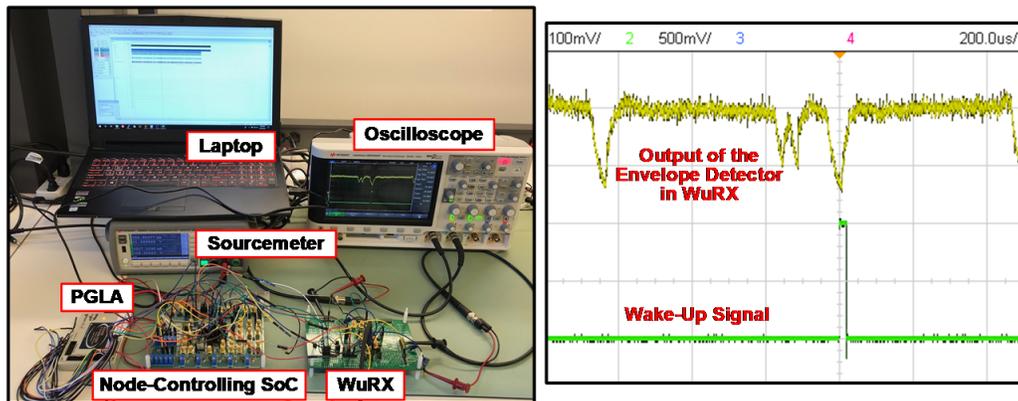


Fig. 5.14. The combined test setup with the SoC and WuRX, and the measured wakeup signal.

5.4 Conclusion

The contributions presented in this chapter demonstrate the technology readiness of ULP and self-powered systems to perform commercially relevant applications. The health-monitoring SoC demonstrates free fall detection, temperature sensing, and arrhythmia detection, all at ultra-low power levels. It also demonstrates vigilant cardiac monitoring powered entirely by body heat, in a compact wearable form that enables comfortable long-term use. The wireless wake-up and control system demonstrates MEL power mode management and standby energy elimination, a new application with high commercial value that is not possible with existing high-power commercially available circuits. Ultimately, these contributions in system and application design will help accelerate the proliferation of IoT devices into the commercial space.

Chapter 6

Conclusion

The IoT has the potential to radically improve human life over the next few decades by producing tremendous amounts of information. This information will create new insights, leading to better decision making, opportunities for optimization, cost reduction, and improved outcomes. The proliferation of IoT sensing devices is currently being hindered by the challenge of large-scale battery recharge and replacement though, which increases the cost of device deployment. Ultimately, the contributions presented in this dissertation aim address this and accelerate the proliferation of IoT sensing devices.

6.1 Summary of Contributions

Chapter 2: Ultra-Low Power Sensor Data Compression

- The performance of the compression algorithm GAS-LEC was evaluated on biomedical sensor data (prior art only considered environmental sensor data). It achieved an average CR of 2.4 for ECG data from the MIT-BIH Arrhythmia Database [46], and 2.0 for acceleration data from the UCI Machine Learning Repository [49].
- GAS-LEC was implemented on a health-monitoring SoC as a hardware accelerator, and fabricated in a 130 nm CMOS technology.
- The accelerator was integrated with the TX interface, realizing a custom data-flow architecture [33] where it can be used without active processor intervention, minimizing its contribution to SoC power.

- The accelerator adds only 4.4 nW (32 kHz, 0.5 V) processing power overhead, and achieves a CR of 2.39, the lowest power and highest CR reported at the time of publication for lossless on-chip ECG data compression.
- The accelerator reduces the required transmitter duty cycle by 3.7x, reducing the system power by 2.9x, and allows the entire system to consume just 2.62 μ W when transmitting ECG data at a 360 Hz sampling rate.

Chapter 3: Ultra-Low Power DVFS RISC-V Microprocessor

- A RISC-V microprocessor core was implemented using a novel SDLS logic style, and fabricated in a 65 nm low-power CMOS process.
- Together with the ACG and VSC, the core realizes a fully integrated modified DVFS scheme that enables nW-level performance flexibility for self-powered IoT sensor nodes in energy-scarce environments.
- At a constant 0.6 V supply voltage, the core can scale its performance from 6 nW at 11 Hz to 140 nW at 8.2 kHz.
- Across the supply voltage range, the core realizes a minimum power of 840 pW, a maximum frequency of 41.2 kHz, and a minimum energy of 13.4 pJ/cycle.
- This work addresses the poor scaling granularity of prior art [35], and unlocks a previously inaccessible area of the IoT application space, shown Fig. 3.8.
- The achievable Hz-kHz frequency range is ideal for many low-frequency IoT sensing applications.

Chapter 4: Multi-Threshold DLS Logic

- A gate-level multi-threshold technique for power reduction in DLS logic in 65 nm CMOS was presented.

- MT-DLS was shown to effectively bridge the power-performance between the two different variants of DLS logic gates (L_{VT} FBB and IO FBB).
- An array of 8-bit MAC blocks was implemented using MT-DLS, and designed to target different trade-offs between power and performance. MAC blocks were also implemented using H_{VT} static CMOS and single-threshold DLS logic variants for comparison.
- For a 0.4 V supply voltage, MT-DLS enables power savings of up to 6.5x (85%), and power savings throughout the approximate approximate 10 Hz–200 Hz range.
- MT-DLS yields power savings in the 10 Hz–1.4 kHz range across the full 0.2 V–1.0 V supply voltage range, as shown in Fig. 4.8.

Chapter 5: Ultra-Low Power IoT Systems and Applications

- A health-monitoring ULP SoC fabricated in 130 nm CMOS was presented, along with its applications, and integration into a wearable self-powered cardiac monitoring system.
- The SoC consumes 507 nW for the presented free fall detection application, 519 μ W for the temperature sensing, and 1.02 μ W for the arrhythmia detection.
- The SoC was integrated with custom components including a BLE compliant TX, an array of flexible TEGs, a flexible antenna, low-leakage super-capacitor, and an e-textile shirt with dry electrodes. Together, these components realize a vigilant cardiac monitoring system for comfortable long-term use.
- The wearable cardiac monitoring system is powered entirely by body heat and consumes just 65 μ W.
- A second ULP SoC was presented as part of a wireless wake-up and control system for managing the power modes of MELs, and eliminating their phantom standby energy consumption.

- The SoC was fabricated in a low-power 65 nm CMOS technology, and consumes 85 nW for the presented MEL control application.

6.2 Conclusions and Open Problems

This dissertation presents contributions in integrated circuit and self-powered system design for reducing the power consumption of IoT sensing devices, and accelerating their adoption. Digital circuits, digital circuit design techniques, and system-level strategies are presented for reducing the power consumption of wireless TX and digital processing circuits in ULP SoCs. This dissertation also presents two ULP SoCs, their applications, and integration into IoT sensing systems for the purpose of serving as proof of technology readiness, and encouraging commercial adoption. The following sections review some of the key insights presented in each chapter, and also discuss open problems that could be addressed in future research.

Chapter 2 explores the design space for on-chip sensor data compression used for the purpose of minimizing wireless TX power in ULP SoCs. The following principles should be adhered to when selecting a data compression algorithm for this purpose. First, the algorithm should be computationally efficient to minimize the processing power overhead of compression. Second, the algorithm should maximize the CR to minimize the amount of data that must be transmitted or stored. And third, the algorithm should account for the specific sensor data, since the CR of a compression algorithm is highly dependent on the characteristics of the data being compressed. System power reduction can be ensured if the following condition is satisfied:

$$P_{\text{comp}} + \frac{P_{\text{TX}}}{\text{CR}} \leq P_{\text{TX}} \quad (8)$$

P_{comp} represents the additional power consumed by compression, and P_{TX} is the TX power when actively transmitting data. Note that this equation makes the assumption that the standby power consumption of the TX is insignificant compared to the active power.

At this time, algorithms that use only integer-based operations [41–45, 48] offer the best

trade-off between CR and processing power overhead. If compression is implemented on an SoC, it should be integrated directly on the TX data path as done in this work. Otherwise, the microprocessor core must waste clock cycles moving data in and out of the compression accelerator, which also consumes more instruction memory, and potentially data memory to store the compressed data. In summary, the results presented in this chapter demonstrate that on-chip sensor data compression can be implemented with very low processing power overhead, and therefore that it is an effective strategy for power reduction in wireless SoCs. It is also widely-applicable, as algorithms such as GAS-LEC are effective for any sensor data type where there is a significant temporal correlation between consecutive samples. It is useful for any application where sensor data must be wirelessly transmitted, and most effective for cases where the TX already dominates the overall system power consumption. It is also completely agnostic to the TX circuit design itself, and therefore a complementary approach that can stack with future improvements in TX circuit design.

One open problem in this design space is that algorithms such as GAS-LEC can compound data loss when there is packet loss between the TX and Receiver (RX). These algorithms rely on encoding and transmitting the differences between consecutive samples, and an actual sample must initially be sent to orientate the decoding process. If a packet is lost, the data that comes after the lost packet cannot be decoded, as shown in Fig. 6.1. One potential solution to this problem is to periodically transmit an actual sample. The max data loss then becomes limited by this period, but this adversely effects the CR, so it is important to maximize this period while accounting for the maximum amount of data loss that can be tolerated by the application. Future work could investigate this trade-off more for different packet loss rates and application requirements.

Chapters 3 and 4 explore new methods for implementing and optimizing low-frequency digital circuits that can operate at lower power levels than those attainable with conventional static CMOS logic. As the clock frequency is scaled down, static CMOS circuit become limited by static leakage currents during active-mode operation. Although standby modes realized with

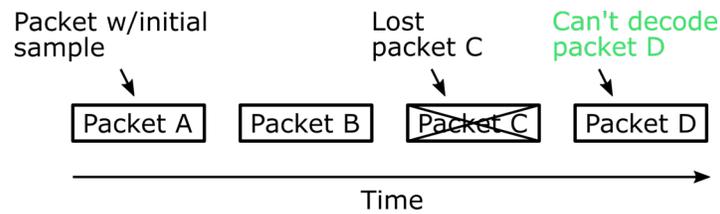


Fig. 6.1. Packet loss between the TX and RX can prevent the data from being decoded correctly.

cutoff structures can bypass this limit, they imply intermittent operation and introduce overheads associated with mode transitions and energy storage. Prior art [34, 59] introduced DLS logic for facilitating active-mode operation at power levels below the leakage floor of static CMOS, but its usefulness was limited by the extremely low clock frequency. Further work [35] added bypass transistors to the DLS logic gate to enable dual-mode operation, which allows circuits to be dynamically reconfigured in either a DLS mode of operation, or a static CMOS mode of operation. The power-performance difference between these two modes of operation is large though, and inflexible for circuits operating directly from harvested energy without a significant energy storage buffer.

Chapter 3 of this dissertation demonstrates that by actively biasing the gate voltage of the bypass transistors (M_{CN} and M_{CP}) in the DLS logic gate, a continuous power-performance trade-off can be realized. This is shown in Fig. 3.8. This is currently the best method for realizing a dynamic power-performance trade-off in DLS logic, as conventional voltage scaling does not provide a significant trade-off, since the on-state current of DLS logic gates is subthreshold leakage, which is weakly dependant on supply voltage. Compared to static CMOS logic implemented with H_{VT} devices, SDLS logic consumes significantly more power when operating in its fastest mode. This is because SDLS logic gates have worse dynamic power consumption due to increased logic gate output capacitance (due to M_{HN} and M_{FP}). SDLS logic can reach much lower power levels in its slower operating modes though. It is therefore best suited for applications where circuits spend most of their time operating at low Hz range frequencies, but need occasional kHz range performance. It is particularly well suited for size-constrained applications without significant energy storage, wherein power must be consumed

as it is harvested. The sub-nW power floor means circuits can actively operate in poor energy harvesting conditions with a very small energy harvester. SDLS logic unfortunately requires significantly more area than static CMOS. The SDLS inverter logic gate is approximately 4x larger than the static CMOS inverter, as previously shown in Fig. 3.9. This could hinder large-scale commercial adoption due to the associated cost increase compared to static CMOS.

Chapter 4 of this dissertation explores a gate-level multi-threshold technique for power reduction in DLS logic. Prior art [61] [JB15] demonstrated several variants of FBB DLS logic in 65 nm, which are preferable over SDLS due to their lower implementation overhead for applications where the capability to dynamically trade-off power and performance isn't useful. These variants have vastly different achievable power and performance ranges though, as previously shown in Fig. 4.2. The results presented in this dissertation demonstrate that these variants can be effectively combined within the same circuit to realize power savings. The faster but higher power L_{VT} FBB logic gates can be selectively used on timing critical paths as necessary to satisfy the application frequency constraints, and the slower but lower power IO FBB DLS gates can be used otherwise. The magnitude of the power savings depends significantly on clock frequency constraint required by the application, and also the specific circuit being implemented. For the 8-bit MAC block studied in this work, power savings are realized in the approximate 10 Hz–1.4 kHz range. The results also show that the leakage power consumption of MT-DLS circuits is approximately linear with the percentage of L_{VT} FBB gates, since they are much higher power than the IO FBB gates. The relationship between the max frequency and percentage of L_{VT} FBB gates is circuit specific though, and depends on the timing path distribution. Future work could explore this relationship for other circuits, as larger circuits with more non timing critical paths could benefit significantly more from MT-DLS than the 8-bit MAC block studied in this work.

Chapter 5 first presents a self-powered health-monitoring SoC and explores its application space. Free fall detection, temperature sensing, and arrhythmia detection applications are developed and implemented, demonstrating the capability of self-powered systems to perform

commercially relevant applications. The results demonstrate that existing commercial products, such as those used for impact tracking in shipping and handling [62] could be replaced with a battery-less solution, reducing their maintenance costs. A fully-wearable cardiac monitoring system is also presented, demonstrating that reliable self-powered operation from body heat is achievable in a comfortable form. Future work could easily make this system more like a commercial product by increasing the levels of integration. The TX and the DC-DC converter used for the TEGs could be integrated with the SoC onto a single IC, the three PCBs could be replaced with one, and the PCB and its case could be further optimized to minimize their size. Several bug fixes could also be implemented for power management circuits on the SoC, which would likely reduce the system power further by 10s of μW s, enabling smaller TEGs to be used.

6.3 Individual Contributions

Many of the projects I've worked on in graduate school are large and complex, and thus multiple people contributed them. This is an explicit list of my contributions for all of the content in this dissertation.

Chapter 2: Ultra-Low Power Sensor Data Compression

All of the content presented in this section was my own except for the physical RTL-to-GDS implementation of the accelerator, which was done by Christopher J. Lukas and Farah Yahya, since the accelerator was integrated into a larger block on the health-monitoring SoC.

Chapter 3: Ultra-Low Power DVFS RISC-V Microprocessor

My contributions focus on the circuit design with the SDLS logic gates, and include the logic gate characterization, digital subsystem design, and digital subsystem implementation. The SDLS logic gates, VSC, and ACG were designed by Daniel S. Truesdell. The SRAM was designed by Sumanth Kamineni and Ningxi Liu.

Chapter 4: Multi-Threshold DLS Logic

All of the content presented in this section was my own except for the design of the DLS

logic gates, which were designed by Daniel S. Truesdell.

Chapter 5: Ultra-Low Power IoT Systems and Applications

The health-monitoring SoC was designed by a large team led by Farah Yahya and Christopher J. Lukas. My contributions to the SoC design were the sensor data compression accelerator, and digital subsystem verification during the design stage. I also designed and implemented the applications used to demonstrate the SoC, and measured the results. I co-lead the integration efforts for the self-powered wearable cardiac monitoring system with Luis Lopez Ruiz. I designed and implemented the digital subsystem for the MEL control SoC, and the digital back-end for the 802.11ba WuRX.

Appendix A

Publications

A.1 Completed

- [JB1] Banerjee, A., **J. Breiholz**, and B.H. Calhoun, "A 130nm Canary SRAM for SRAM Dynamic Write VMIN Tracking across Voltage, Frequency, and Temperature Variations," in 2015 IEEE Custom Integrated Circuits Conference (CICC).
- [JB2] **J. Breiholz**, F. Yahya, C.J. Lukas, X. Chen, K. Leach, D. Wentzloff, and B.H. Calhoun. "A 4.4 nW lossless sensor data compression accelerator for 2.9x system power reduction in wireless body sensors," in 2017 IEEE Midwest Symposium on Circuits and Systems (MWSCAS).
- [JB3] F. Yahya, C. Lukas, **J. Breiholz**, A. Roy, H.N. Patel, N. Liu, X. Chen, A. Kosari, S. Li, D. Akella, O. Ayorinde, D. Wentzloff, and B.H. Calhoun, "A battery-less 507nW SoC with integrated platform power manager and SiP interfaces," in 2017 IEEE Symposium on VLSI Circuits (VLSI).
- [JB4] A. Kosari, **J. Breiholz**, N. Liu, B.H. Calhoun, and D.D. Wentzloff, "A 0.5 V 68 nW ECG Monitoring Analog Front-End for Arrhythmia Diagnosis" in 2018 Journal of Low Power Electronics and Applications (JLPEA).
- [JB5] A. Alghaihab, **J. Breiholz**, H. Kim, B. Calhoun, and D. D. Wentzloff, "A 150 μ W-57.5 dBm-Sensitivity Bluetooth Low-Energy Back-Channel Receiver with LO Frequency Hopping," in 2018 IEEE Radio Frequency Integrated Circuits Symposium (RFIC).
- [JB6] X. Chen, **J. Breiholz**, F. Yahya, C. Lukas, H.S. Kim, B.H. Calhoun, and D. Wentzloff. "A 486 μ W All-Digital Bluetooth Low Energy Transmitter with Ring Oscillator Based ADPLL for IoT applications," in 2018 IEEE Radio Frequency Integrated Circuits Symposium (RFIC).
- [JB7] C. Lukas, F. Yahya, **J. Breiholz**, A. Roy, X. Chen, H.N. Patel, N. Liu, A. Kosari, S. Li, D. Akella, O. Ayorinde, D.D. Wentzloff, and B.H. Calhoun. "A 1.02 μ W Battery-less, Continuous Sensing and Post-processing SiP for Wearable Applications," in 2018 IEEE Transactions on Biomedical Circuits and Systems (TBCAS)
- [JB8] A. Alghaihab, Y. Shi, **J. Breiholz**, H.K. Kim, B.H. Calhoun, and D.D. Wentzloff. "Enhanced Interference Rejection Bluetooth Low-Energy Back-Channel Receiver with LO Frequency Hopping," in 2019 IEEE Journal of Solid-State Circuits (JSSC).

- [JB9] X. Chen, **J. Breiholz**, F. Yahya, C.J. Lukas, H.K. Kim, B.H. Calhoun, and D.D. Wentzloff. "Analysis and Design of an Ultra-Low-Power Bluetooth Low Energy Transmitter with Ring Oscillator Based ADPLL and 4X Frequency Edge Combiner," in 2019 IEEE Journal of Solid-State Circuits (JSSC).
- [JB10] D. Truesdell, **J. Breiholz**, S. Kamineni, N. Liu, A. Magyar, B.H. Calhoun. "A 5.8nW-86.5nW 6Hz-4.4kHz DVFS RISC-V Microprocessor Using Scalable Dynamic Leakage Suppression Logic for Battery-less IoT Applications," in 2019 IEEE Solid-State Circuits Letters (LSSC).
- [JB11] S. Li, **J. Breiholz**, S. Kamineni, J. Im, D.D. Wentzloff, and B.H. Calhoun. "An 85 nW IoT Node-Controlling SoC for MELs Power-Mode Management and Phantom Energy Reduction," in 2020 IEEE International Symposium on Circuits and Systems (ISCAS). **Best Student Paper Award Winner.**
- [JB12] J. Im, **J. Breiholz**, S. Li, B.H. Calhoun, and D.D. Wentzloff. "A Fully Integrated 0.2 V 802.11ba Wake-Up Receiver with -91.5 dBm Sensitivity," in 2020 IEEE Radio Frequency Integrated Circuits Symposium (RFIC).

A.2 Planned

- [JB13] Paper on Multi-Threshold Technique for Dynamic Leakage Suppression Logic
- [JB14] Journal paper on wearable cardiac monitoring system power by body heat. Planned for Nature, draft is complete and pending review by coauthors.
- [JB15] Journal paper on DLS logic. Planned for IEEE Journal of Solid-State Circuits (JSSC), work in progress.
- [JB16] Conference paper on next generation SoC for MEL control. Pending tape-out re-spin for bug fixes, work in progress.
- [JB17] Journal paper on 802.11ba Wake-Up Receiver. Planned for IEEE Transactions on Microwave Theory and Techniques (TMTT), submitted.

Acronyms

ACG Adaptive Clock Generator. 35, 39, 42, 43, 69, 75

ADC Analog to Digital Converter. x, 55, 58, 60

AFE Analog Front End. x, 9, 53, 58

AFib Atrial Fibrillation. x, 9, 13, 59, 61

APR Automated Place and Route. 49

BLE Bluetooth Low Energy. xi, 59, 60, 70

CMOS Complementary Metal-Oxide-Semiconductor. iv, ix, x, 6–9, 16–18, 26, 27, 29–37, 39, 43, 45–49, 51, 54, 64, 68–74

CR Compression Ratio. xii, 8, 12, 14–18, 23–25, 68, 69, 71, 72

DIBL Drain-Induced Barrier Lowering. 47

DLS Dynamic Leakage Suppression. iv, v, ix, x, 8, 9, 32–36, 39, 40, 46–49, 51, 52, 70, 73–75

DoE Department of Energy. 62

DSP Digital Signal Processing. 48, 54

DVFS Dynamic Voltage and Frequency Scaling. x, 27, 32, 36, 39–44, 69

ECG Electrocardiogram. iv, ix–xii, 5, 6, 8, 13, 16, 18, 20–25, 58–61, 68, 69

FA-LEC Frequency-Adaptive Lossless Entropy Compression. 17

FAS-LEC Frequency-Adaptive Split Lossless Entropy Compression. 17, 18

FBB Forward Body Bias. x, 46–52, 70, 74

FFT Fast Fourier Transform. 48

FIFO First-In First-Out. 56

FSK Frequency Shift Keying. 23

GA-LEC Greedy-Adaptive Lossless Entropy Compression. 17, 23

GAS-LEC Greedy-Adaptive Split Lossless Entropy Compression. ix, 17–19, 21, 23, 24, 68, 72

- GDS** Graphic Data System. 75
- GIDL** Gate-Induced Drain Leakage. 46
- GIF** Graphics Interchange Format. 14
- GPIO** General-Purpose Input Output. x, xi, 39–41, 56, 64–66
- IC** Integrated Circuit. 4, 75
- IoT** Internet of Things. iii, iv, ix, xii, 1–9, 11–14, 17, 25–27, 29, 31, 32, 36, 43, 45, 46, 51, 53–55, 57, 67–69, 71
- LDO** Low-Dropout. 64
- LEC** Lossless Entropy Compression. ix, 15–17, 23
- LED** Light-Emitting Diode. 62
- LZW** Lempel-Ziv-Welch. 14, 23
- MAC** Multiply-Accumulate. iv, x, 9, 48–52, 70, 74
- MEL** Miscellaneous Electric Load. xi, 9, 53, 61–67, 70, 71, 76
- MEMS** Micro-Electro-Mechanical System. 55, 57
- MIT-BIH** Massachusetts Institute of Technology Beth Israel Hospital. ix, 16, 17, 20–22, 24, 25, 68
- MT-DLS** Multi-Threshold Dynamic Leakage Suppression. x, 9, 47–52, 70, 74
- PCB** Printed Circuit Board. xi, 59, 60, 75
- PMU** Power Management Unit. 54
- PPG** Photoplethysmography. 5
- PV** Photovoltaic. 3, 54
- RAM** Random-Access Memory. 15
- RBB** Reverse Body Biasing. 28
- RF** Radio Frequency. 2, 4
- ROM** Read-Only Memory. 15
- RTL** Register-Transfer Level. 75
- RX** Receiver. xi, 72, 73

- S-LZW** Sensor Lempel-Ziv-Welch. 14–16
- SAR** Successive Approximation Register. 58
- SDLS** Scalable Dynamic Leakage Suppression. ix, x, 8, 9, 35–46, 69, 73–75
- SoC** System on Chip. iii–v, ix–xi, 4–10, 13–15, 18, 22–25, 31, 53–60, 63–68, 70–72, 74–76
- SPI** Serial Peripheral Interface. 39, 55, 56, 64, 66
- SRAM** Static Random-Access Memory. 39, 54, 64, 75
- Sub-V_t** Sub-Threshold. 49, 50
- TEG** Thermoelectric Generator. 3, 54, 60, 61, 70, 75
- TT** Typical Typical. 49
- TX** Transmitter. iii, iv, ix, xi, 5–8, 13, 14, 18, 23–25, 53, 55–61, 68, 70–73, 75
- UCI** University of California, Irvine. ix, 20, 22, 24, 25, 68
- ULP** Ultra-Low Power. iii, iv, 4–9, 13–15, 18, 25, 53, 63, 67, 70, 71
- VSC** Voltage Scaling Controller. x, 36, 39, 40, 42, 43, 69, 75
- VTC** Voltage Transfer Characteristic. ix, 34
- WuRX** Wake-Up Receiver. xi, 9, 54, 63–67, 76

References

- [1] P. Sparks, "The route to a trillion devices: The outlook for iot investment to 2035," Arm, Tech. Rep., June 2017. [Online]. Available: <https://learn.arm.com/route-to-trillion-devices.html>
- [2] *ULP meets energy harvesting: a game-changing combination for design engineers*, Texas Instruments, 2010. [Online]. Available: <https://www.mouser.ph/pdfDocs/TI-ULP-meets-energy-harvesting-A-game-changing-combination-for-design-engineers.pdf>
- [3] Y. Du, J. Xu, B. Paul, and P. Eklund, "Flexible thermoelectric materials and devices," *Applied Materials Today*, vol. 12, pp. 366 – 388, 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2352940718302804>
- [4] M. Piñuela, P. D. Mitcheson, and S. Lucyszyn, "Ambient rf energy harvesting in urban and semi-urban environments," *IEEE Transactions on Microwave Theory and Techniques*, vol. 61, no. 7, pp. 2715–2726, 2013.
- [5] X. Liu, H. Gao, J. E. Ward, X. Liu, B. Yin, T. Fu, J. Chen, D. R. Lovley, and J. Yao, "Power generation from ambient humidity using protein nanowires," *Nature*, vol. 578, no. 7796, pp. 550–554, Feb 2020. [Online]. Available: <https://doi.org/10.1038/s41586-020-2010-9>
- [6] B. J. Hansen, Y. Liu, R. Yang, and Z. L. Wang, "Hybrid nanogenerator for concurrently harvesting biomechanical and biochemical energy," *ACS Nano*, vol. 4, no. 7, pp. 3647–3652, Jul 2010. [Online]. Available: <https://doi.org/10.1021/nn100845b>
- [7] X. Wu, I. Lee, Q. Dong, K. Yang, D. Kim, J. Wang, Y. Peng, Y. Zhang, M. Saligane, M. Yasuda, K. Kumeno, F. Ohno, S. Miyoshi, M. Kawaminami, D. Sylvester, and D. Blaauw, "A 0.04mm² 316nm wireless and batteryless sensor system with integrated cortex-m0+ processor and optical communication for cellular temperature measurement," in *2018 IEEE Symposium on VLSI Circuits*, 2018, pp. 191–192.
- [8] S. Han, J. Kim, S. M. Won, Y. Ma, D. Kang, Z. Xie, K.-T. Lee, H. U. Chung, A. Banks, S. Min, S. Y. Heo, C. R. Davies, J. W. Lee, C.-H. Lee, B. H. Kim, K. Li, Y. Zhou, C. Wei, X. Feng, Y. Huang, and J. A. Rogers, "Battery-free, wireless sensors for full-body pressure and temperature mapping," *Science Translational Medicine*, vol. 10, no. 435, 2018. [Online]. Available: <https://stm.sciencemag.org/content/10/435/eaan4950>
- [9] Y. Zhang, F. Zhang, Y. Shakhsher, J. D. Silver, A. Klinefelter, M. Nagaraju, J. Boley, J. Pandey, A. Shrivastava, E. J. Carlson, A. Wood, B. H. Calhoun, and B. P. Otis, "A batteryless 19 μ w mics/ism-band energy harvesting body sensor node soc for exg applications," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 1, pp. 199–213, Jan 2013.

- [10] A. Roy, A. Klinefelter, F. B. Yahya, X. Chen, L. P. Gonzalez-Guerrero, C. J. Lukas, D. A. Kamakshi, J. Boley, K. Craig, M. Faisal, S. Oh, N. E. Roberts, Y. Shakhsheer, A. Shrivastava, D. P. Vasudevan, D. D. Wentzloff, and B. H. Calhoun, "A 6.45 μ w self-powered soc with integrated energy-harvesting power management and ulp asymmetric radios for portable biomedical systems," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 9, no. 6, pp. 862–874, Dec 2015.
- [11] J. Myers, A. Savanth, R. Gaddh, D. Howard, P. Prabhat, and D. Flynn, "A subthreshold arm cortex-m0+ subsystem in 65 nm cmos for wsn applications with 14 power domains, 10t sram, and integrated voltage regulator," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 1, pp. 31–44, Jan 2016.
- [12] J. K. Brown, D. Abdallah, J. Boley, N. Collins, K. Craig, G. Glennon, K. Huang, C. J. Lukas, W. Moore, R. K. Sawyer, Y. Shakhsheer, F. B. Yahya, A. Wang, N. E. Roberts, D. D. Wentzloff, and B. H. Calhoun, "A 65nm energy-harvesting ulp soc with 256kb cortex-m0 enabling an 89.1 μ w continuous machine health monitoring wireless self-powered system," in *2020 IEEE International Solid-State Circuits Conference - (ISSCC)*, 2020, pp. 420–422.
- [13] *STM32L0 series of ultra-low-power MCUs*, STMicroelectronics, 2020. [Online]. Available: <https://www.st.com/en/microcontrollers-microprocessors/stm32l0-series.html>
- [14] H. Kim, S. Kim, N. Van Helleputte, A. Artes, M. Konijnenburg, J. Huisken, C. Van Hoof, and R. F. Yazicioglu, "A configurable and low-power mixed signal soc for portable ecg monitoring applications," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 8, no. 2, pp. 257–267, 2014.
- [15] F. B. Yahya, C. J. Lukas, and B. H. Calhoun, "A top-down approach to building battery-less self-powered systems for the internet-of-things," *Journal of Low Power Electronics and Applications*, vol. 8, no. 2, 2018. [Online]. Available: <https://www.mdpi.com/2079-9268/8/2/21>
- [16] O. Kwon, J. Jeong, H. Kim, I. Kwon, S. Park, J. Kim, and Y. Choi, "Electrocardiogram sampling frequency range acceptable for heart rate variability analysis," in *Healthcare informatics research*, vol. 24(3), 2018, pp. 198–206.
- [17] A. Sharma, Seung Bae Lee, A. Polley, S. Narayanan, Wen Li, T. Sculley, and S. Ramaswamy, "Multi-modal smart bio-sensing soc platform with >80db snr 35a ppg rx chain," in *2016 IEEE Symposium on VLSI Circuits (VLSI-Circuits)*, 2016, pp. 1–2.
- [18] Song-Mi Lee, Sang Min Yoon, and Heeryon Cho, "Human activity recognition from accelerometer data using convolutional neural network," in *2017 IEEE International Conference on Big Data and Smart Computing (BigComp)*, 2017, pp. 131–134.
- [19] P. Casale, O. Pujol, and P. Radeva, "Human activity recognition from accelerometer data using a wearable device," vol. 6669, 06 2011, pp. 289–296.
- [20] C. N. Teague, J. A. Heller, B. N. Nevius, A. M. Carek, S. Mabrouk, F. Garcia-Vicente, O. T. Inan, and M. Etemadi, "A wearable, multimodal sensing system to monitor knee joint health," *IEEE Sensors Journal*, vol. 20, no. 18, pp. 10 323–10 334, 2020.

- [21] W. Gao, S. Emaminejad, H. Y. Y. Nyein, S. Challa, K. Chen, A. Peck, H. M. Fahad, H. Ota, H. Shiraki, D. Kiriya, D.-H. Lien, G. A. Brooks, R. W. Davis, and A. Javey, "Fully integrated wearable sensor arrays for multiplexed in situ perspiration analysis," *Nature*, vol. 529, no. 7587, pp. 509–514, Jan 2016. [Online]. Available: <https://doi.org/10.1038/nature16521>
- [22] M. Moghaddam, A. Silva, D. Clewley, R. Akbar, S. Hussaini, J. Whitcomb, R. Devarakonda, R. Shrestha, R. Cook, G. Prakash, S. Santhana Vannan, and A. Boyer. (2017) Soil moisture profiles and temperature data from soilscape sites, usa. ornl daac, oak ridge, tennessee, usa. [Online]. Available: <https://doi.org/10.3334/ORNLDAAAC/1339>
- [23] V. N. Deekshit, M. V. Ramesh, P. K. Indukala, and G. J. Nair, "Smart geophone sensor network for effective detection of landslide induced geophone signals," in *2016 International Conference on Communication and Signal Processing (ICCSP)*, 2016, pp. 1565–1569.
- [24] M. Sophocleous, P. Savva, M. F. Petrou, J. K. Atkinson, and J. Georgiou, "A durable, screen-printed sensor for in situ and real-time monitoring of concrete's electrical resistivity suitable for smart buildings/cities and iot," *IEEE Sensors Letters*, vol. 2, no. 4, pp. 1–4, 2018.
- [25] W. Tushar, N. Wijerathne, W. Li, C. Yuen, H. V. Poor, T. K. Saha, and K. L. Wood, "Internet of things for green building management: Disruptive innovations through low-cost sensor technology and artificial intelligence," *IEEE Signal Processing Magazine*, vol. 35, no. 5, pp. 100–110, 2018.
- [26] C. Chiang, Y. Lu, and L. Lin, "A cmos fish spoilage detector for iot applications of fish markets," *IEEE Sensors Journal*, vol. 18, no. 1, pp. 375–381, 2018.
- [27] F. Wanlass and C. Sah, "Nanowatt logic using field-effect metal-oxide semiconductor triodes," in *1963 IEEE International Solid-State Circuits Conference. Digest of Technical Papers*, vol. VI, 1963, pp. 32–33.
- [28] S. Hanson, M. Seok, Y. Lin, Z. Foo, D. Kim, Y. Lee, N. Liu, D. Sylvester, and D. Blaauw, "A low-voltage processor for sensing applications with picowatt standby mode," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 4, pp. 1145–1155, 2009.
- [29] F. Yahya, C. J. Lukas, J. Breiholz, A. Roy, H. N. Patel, N. Liu, X. Chen, A. Kosari, S. Li, D. Akella, O. Ayorinde, D. Wentzloff, and B. H. Calhoun, "A battery-less 507nw soc with integrated platform power manager and sip interfaces," in *2017 Symposium on VLSI Circuits*, June 2017, pp. C338–C339.
- [30] Everactive. (2020). [Online]. Available: <https://everactive.com/>
- [31] e-peas Semiconductors. (2020). [Online]. Available: <https://e-peas.com/>
- [32] S. Inc. (2020). [Online]. Available: <https://www.steamiq.com/>
- [33] C. Lukas, B. Calhoun, and F. Yahya, "System on a chip with customized data flow architecture," 2017, patent US 2020/0065278 A1. [Online]. Available: <https://patents.google.com/patent/US20200065278A1/en>

- [34] W. Lim, I. Lee, D. Sylvester, and D. Blaauw, "8.2 batteryless sub-nw cortex-m0+ processor with dynamic leakage-suppression logic," in *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, Feb 2015, pp. 1–3.
- [35] L. Lin, S. Jain, and M. Alioto, "A 595pw 14pj/cycle microcontroller with dual-mode standard cells and self-startup for battery-indifferent distributed sensing," in *2018 IEEE International Solid - State Circuits Conference - (ISSCC)*, Feb 2018, pp. 44–46.
- [36] M. L. Hilton, "Wavelet and wavelet packet compression of electrocardiograms," *IEEE Transactions on Biomedical Engineering*, vol. 44, no. 5, pp. 394–402, 1997.
- [37] F. Zhang, Y. Zhang, J. Silver, Y. Shakhsher, M. Nagaraju, A. Klinefelter, J. Pandey, J. Boley, E. Carlson, A. Shrivastava, B. Otis, and B. Calhoun, "A batteryless 19 μ w mics/ism-band energy harvesting body area sensor node soc," in *2012 IEEE International Solid-State Circuits Conference*, 2012, pp. 298–300.
- [38] S. J. Raudys and A. K. Jain, "Small sample size effects in statistical pattern recognition: recommendations for practitioners," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 3, pp. 252–264, 1991.
- [39] Wikipedia. (2020) Lempel–ziv–welch. [Online]. Available: <https://en.wikipedia.org/wiki/Lempel-Ziv-Welch>
- [40] C. Sadler and M. Martonosi, "Data compression algorithms for energy-constrained devices in delay tolerant networks," 10 2006, pp. 265–278.
- [41] F. Marcelloni and M. Vecchio, "A simple algorithm for data compression in wireless sensor networks," *IEEE Communications Letters*, vol. 12, no. 6, pp. 411–413, 2008.
- [42] E. Chua and W. C. Fang, "Mixed bio-signal lossless data compressor for portable brain-heart monitoring systems," *IEEE Transactions on Consumer Electronics*, vol. 57, no. 1, pp. 267–273, February 2011.
- [43] C. J. Deepu, X. Zhang, W. S. Liew, D. L. T. Wong, and Y. Lian, "An ecg-soc with 535nw/channel lossless data compression for wearable sensors," in *2013 IEEE Asian Solid-State Circuits Conference (A-SSCC)*, 2013, pp. 145–148.
- [44] Y. Liang and Y. Li, "An efficient and robust data compression algorithm in wireless sensor networks," *IEEE Communications Letters*, vol. 18, no. 3, pp. 439–442, 2014.
- [45] D. Yunge, S. Park, P. Kindt, and S. Chakraborty, "Dynamic alternation of huffman code-books for sensor data compression," *IEEE Embedded Systems Letters*, vol. 9, no. 3, pp. 81–84, 2017.
- [46] MIT-BIH arrhythmia database. [Online]. Available: <https://physionet.org/physiobank/database/mitdb/>
- [47] C. J. Deepu, X. Zhang, W. S. Liew, D. L. T. Wong, and Y. Lian, "An ecg-on-chip with 535 nw/channel integrated lossless data compressor for wireless sensors," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 11, pp. 2435–2448, Nov 2014.

- [48] M. Vecchio, R. Giaffreda, and F. Marcelloni, "Adaptive lossless entropy compressors for tiny iot devices," *IEEE Transactions on Wireless Communications*, vol. 13, no. 2, pp. 1088–1100, February 2014.
- [49] UCI Machine Learning Repository: Activity Recognition from Single Chest-Mounted Accelerometer Data Set. [Online]. Available: <https://archive.ics.uci.edu/ml/datasets/Activity+Recognition+from+Single+Chest-Mounted+Accelerometer>
- [50] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits*, 3rd ed. USA: Prentice Hall Press, 2008.
- [51] M. Seok, D. Sylvester, and D. Blaauw, "Optimal technology selection for minimizing energy and variability in low voltage applications," in *Proceeding of the 13th international symposium on Low power electronics and design (ISLPED '08)*, 2008, pp. 9–14.
- [52] S. Hanson, B. Zhai, M. Seok, B. Cline, K. Zhou, M. Singhal, M. Minuth, J. Olson, L. Nazhandali, T. Austin, D. Sylvester, and D. Blaauw, "Exploring variability and performance in a sub-200-mv processor," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 4, pp. 881–891, 2008.
- [53] B. H. Calhoun, F. A. Honore, and A. P. Chandrakasan, "A leakage reduction methodology for distributed mtcmos," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 5, pp. 818–826, 2004.
- [54] B. H. Calhoun and A. P. Chandrakasan, "Ultra-dynamic voltage scaling (udvs) using sub-threshold operation and local voltage dithering," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 1, pp. 238–245, 2006.
- [55] L. Lin, S. Jain, and M. Alioto, "Multi-sensor platform with five-order-of-magnitude system power adaptation down to 3.1nw and sustained operation under moonlight harvesting," in *2020 IEEE Symposium on VLSI Circuits*, 2020, pp. 1–2.
- [56] *Enerchip Bare Die Rechargeable Solid State Bare Die Batteries*, Cymbet Corporation, 2016. [Online]. Available: <https://www.cymbet.com/products/enerchip-solid-state-batteries/>
- [57] *DMHA14R5V353M4ATA0*, Murata Manufacturing Co., 2017. [Online]. Available: <https://www.murata.com/en-us/products/productdetail.aspx?partno=DMHA14R5V353M4ATA0>
- [58] *AMK105EBJ226MV-F*, Taiyo Yuden, 2020. [Online]. Available: <https://www.digikey.com/en/products/detail/taiyo-yuden/AMK105EBJ226MV-F/9916386>
- [59] D. Bol, R. Ambroise, D. Flandre, and J. Legat, "Building ultra-low-power low-frequency digital circuits with high-speed devices," in *2007 14th IEEE International Conference on Electronics, Circuits and Systems*, Dec 2007, pp. 1404–1407.
- [60] K. Asanović, "The rocket chip generator," EECS Department, University of California, Berkeley, Tech. Rep., April 2016. [Online]. Available: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-17.pdf>

- [61] D. S. Truesdell and B. H. Calhoun, "Improving dynamic leakage suppression logic with forward body bias in 65nm cmos," in *2019 IEEE SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S)*, 2019.
- [62] SpotSee. (2020) Impact sensors. [Online]. Available: <https://spotsee.io/impact>
- [63] B. E. Bergen G, Stevens MR, "Falls and fall injuries among adults aged ≥ 65 years — united states, 2014," *MMWR Morb Mortal Wkly Rep* 2016, no. 65, pp. 993–998, 2016.
- [64] *Data Sheet: ADXL362*, Analog Devices, 2019. [Online]. Available: <https://www.analog.com/media/en/technical-documentation/data-sheets/ADXL362.pdf>
- [65] N. Jia, *Detecting Human Fall with a 3-Axis Digital Accelerometer*, Analog Devices, 2009. [Online]. Available: <https://www.analog.com/en/analog-dialogue/articles/detecting-falls-3-axis-digital-accelerometer.html>
- [66] Deloitte. 2017 global health care outlook. [Online]. Available: <https://www2.deloitte.com/us/en/pages/life-sciences-and-health-care/articles/global-health-care-sector-outlook.html>
- [67] "Cardiovascular disease fact sheet," The World Health Organization, Tech. Rep., May 2017. [Online]. Available: [https://www.who.int/en/news-room/fact-sheets/detail/cardiovascular-diseases-\(cvds\)](https://www.who.int/en/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds))
- [68] Wikipedia. (2020) Holter monitor. [Online]. Available: https://en.wikipedia.org/wiki/Holter_monitor
- [69] J. Pan and W. J. Tompkins, "A real-time qrs detection algorithm," *IEEE Transactions on Biomedical Engineering*, vol. BME-32, no. 3, pp. 230–236, 1985.
- [70] D. E. Lake and J. R. Moorman, "Accurate estimation of entropy in very short physiological time series: the problem of atrial fibrillation detection in implanted ventricular devices," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 300, no. 1, pp. H319–H325, 2011, PMID: 21037227. [Online]. Available: <https://doi.org/10.1152/ajpheart.00561.2010>
- [71] D. Fan, L. L. Ruiz, and J. Lach, "Application-driven dynamic power management for self-powered vigilant monitoring," in *2018 IEEE 15th International Conference on Wearable and Implantable Body Sensor Networks (BSN)*, March 2018, pp. 210–213.
- [72] Wikipedia. (2020) Miscellaneous electric load. [Online]. Available: https://en.wikipedia.org/wiki/Miscellaneous_electric_load
- [73] *BTO Investigates Miscellaneous Electric Loads*, U.S. Department of Energy Building Technologies Office, 2016. [Online]. Available: <https://www.energy.gov/eere/buildings/downloads/bto-investigates-miscellaneous-electric-loads>
- [74] W. Wang, J. Su, Z. Hicks, and B. Campbell, "The standby energy of smart devices: Problems, progress, potential," in *2020 IEEE/ACM Fifth International Conference on Internet-of-Things Design and Implementation (IoTDI)*, 2020, pp. 164–175.

- [75] Philips. (2020) Philips hue white a21 e26. [Online]. Available: <https://www.philips-hue.com/en-us/p/hue-white-1-pack-a21-e26/046677557805#specifications>
- [76] D. D. Wentzloff, A. Alghaihab, and J. Im, "Ultra-low power receivers for iot applications: A review," in *2020 IEEE Custom Integrated Circuits Conference (CICC)*, 2020, pp. 1–8.
- [77] T. L. Ericsson. (2020) Wake-up radio - a key component of iot? [Online]. Available: <https://www.ericsson.com/en/blog/2017/12/wake-up-radio--a-key-component-of-iot>
- [78] *MSP430 Ultra-low-power Microcontrollers*, Texas Instruments, 1999. [Online]. Available: <https://www.ti.com/sc/docs/products/micro/msp430/39113.pdf>