

UNIVERSITÀ
DI TRENTO

Department of Industrial Engineering

Doctoral School in Materials, Mechatronics and Systems Engineering

XXXVIII Cycle

Energy Harvesting Circuits for Battery-Free IoT Powered by Plant Microbial Fuel Cells

Supervisor

Prof. Davide Brunelli

Co-supervisors

Prof. Matteo Nardello

Prof. Oscar Casas

Prof. Daniele Fontanelli

PhD Candidate

Maria Doglioni

Copyright Notice

The content of this thesis is the result of the author's original research work. Parts of this thesis are based on material previously published by the author and co-authors in the publications listed below. Text, figures, and tables reproduced or adapted from these publications are included in accordance with the author reuse policies of the respective publishers. The published works are the following:

- M. Doglioni, E. Yildiz, M. Nardello, K. Akhunov, K.S. Yildirim, D. Brunelli, *"CapDYN: Adaptive Self-Scaling Energy Storage for Powering Batteryless IoT,"* ACM Transactions on Embedded Computing Systems, vol. 24, no. 5, Article 77, 2025, pp. 1–32, doi: 10.1145/3737288.
- M. Doglioni, O. Casas, M. Nardello, D. Brunelli, *"Self-powered Plant Microbial Fuel Cell Monitoring Node with Autonomous Dynamic Capacitance,"* Proceedings of the 2025 IEEE International Workshop on Metrology for Industry 4.0 & IoT (MetroInd4.0&IoT), Castelldefels, Spain, July 1–3, 2025, pp. 179-184.
  doi: 10.1109/MetroInd4.0IoT66048.2025.11122107.
- M. Doglioni, R. L. Rosa, M. Nardello, D. Brunelli, *"Energy Harvesting Strategies for Plant Microbial Fuel Cells in Sustainable IoT Applications,"* 2024 IEEE SENSORS, Kobe, Japan, 2024, pp. 1-4.
  doi: 10.1109/SENSORS60989.2024.10784498.
- M. Doglioni, M. Nardello, D. Brunelli, *"Plant Microbial Fuel Cells: Energy Sources and Biosensors for battery-Free Smart Agriculture,"* in IEEE Transactions on AgriFood Electronics,

vol. 2, no. 2, pp. 460-470, Sept.-Oct. 2024.
doi: 10.1109/TAFE.2024.3417644.

- M. Doglioni, A. Torrisi, L. Gemma, A. Albanese, L. Santoro, M. Nardello, D. Brunelli, *"Batteryless Soil EIS Sensor Powered by Microbial Fuel Cell,"* in Proceedings of SIE 2022, G. Cocorullo, F. Crupi, and E. Limiti, Eds., Cham, Switzerland: Springer Nature, 2023, pp. 277–282.
  doi: 10.1007/978-3-031-26066-7_43.
- M. Nardello, M. Doglioni, S. Dagnino, P. Pastorelli, D. Brunelli, *"Every Microjoule Counts: Zero-Failure Task Execution in Batteryless Sensors,"* 2025 IEEE SENSORS, Vancouver, Canada, 2025.
  doi: 10.1109/SENSORS59705.2025.11330908.

For IEEE copyrighted material reused in this thesis, full credit is given to the original publication, followed by the appropriate IEEE copyright notice. Only accepted-manuscript material, and not publisher-formatted IEEE PDFs, is used in the online version of this thesis.

In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of University of Trento's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to `http://www.ieee.org/publications_standards/publications/rights/rights_link.html` to learn how to obtain a License from RightsLink. If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.

Material reproduced or adapted from ACM publications is identified by citation to the corresponding ACM Version of Record and DOI. Such material remains subject to the applicable ACM author reuse and licensing terms.

Material reproduced or adapted from Springer Nature publications is identified by citation to the corresponding published version and remains subject to the applicable Springer Nature author reuse and licensing terms.

Except where otherwise noted, this thesis is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0

International License. Material reproduced or adapted from IEEE, ACM, Springer Nature, or other third-party sources is excluded from this license unless explicitly stated otherwise.

Abstract

Battery-free Internet of Things (IoT) systems can reduce maintenance costs and environmental impact by operating directly from harvested energy. However, replacing batteries with weak and variable sources introduces a fundamental energy-management challenge: harvested power is often in the microwatt range, whereas sensing, computation, communication, and monitoring tasks require short bursts of higher energy. This mismatch is especially critical for Plant Microbial Fuel Cells (PMFCs), which are sustainable bioelectrochemical energy sources but provide low voltage, low power, and slow recovery dynamics.

This dissertation investigates energy-harvesting circuits and storage architectures for battery-free embedded systems powered by PMFCs. First, it presents a controlled intermittent power-extraction interface that imposes explicit extraction and recovery timing while remaining compatible with energy-harvesting converters. Experiments on real PMFCs show that some intermittent extraction periods are associated with higher harvested power: in preliminary tests, a 10 s period reached 83 μW, 30% higher than 100 s and 63% higher than 1000 s. In a 9-day experiment, two intermittently loaded cells increased from 47.2 to 122.4 μW and from 19.4 to 67.6 μW, compared with 32.8 to 43.0 μW for the continuously loaded reference. Electrochemical impedance spectroscopy further shows that loading history affects the anode and cathode responses.

Second, the dissertation presents CapDYN, an autonomous dynamic-capacitance architecture for battery-free intermittent systems. CapDYN adapts the equivalent storage capacitance during charge and discharge, exposes intermediate energy states to the runtime, and matches them to task requirements. The measured hardware reduces leftover energy from 23.67% in a fixed-capacitor baseline to

0.16%–5.57%, while keeping the control infrastructure below 10 μW. In the LTspice-based application-level evaluation, the energy-aware runtime avoids failed executions, improves effective efficiency by up to 34.10%, and advances aggregate first task completions by up to 45.00%.

Overall, this thesis shows that efficient battery-free operation requires co-design across the source interface, storage architecture, and runtime, increasing the fraction of scarce sustainable energy that becomes completed embedded tasks.

Acknowledgements

This thesis would not have been possible without the support and encouragement of many people. First, I would like to thank Prof. Davide Brunelli for his supervision throughout my PhD. His advice and trust have been fundamental to the development of this work.
I am also sincerely grateful to Prof. Oscar Casas and to the ISI group for welcoming me into their laboratory and for creating such an inspiring research environment. Their scientific curiosity and openness gave me the opportunity to broaden my perspective and made my time with them especially valuable.
I would also like to thank Dr. Matteo Nardello, for being a wonderful right arm throughout this work and a rock in difficult moments. His technical support and constant willingness to discuss ideas have been invaluable throughout this work. I am especially grateful for the way he challenged my scientific thinking and helped me grow as a researcher.
I also extend my thanks to Luca Benedetti from the Coatings and Corrosion Lab for his kind assistance with potentiostat operation and electrochemical measurements.
I also thank my colleagues in the Open Space for making my working days lighter, less serious, and certainly less silent. Their company, conversations, and encouragement were especially precious on the days when I needed them most. I am grateful to the colleagues, collaborators, students and friends who shared discussions, breakdowns, experiments, deadlines, coffee breaks, barbecues, bike rides, climbing sessions and the many small moments that made this journey possible and more meaningful.
Finally, I would like to thank my family and loved ones for their patience and support, especially throughout my thesis writing. Grace and Rocky, too.

Οὐκοῦν, ὅταν δὴ μὴ σθένω, πεπαύσομαι.
Ἀρχὴν δὲ θηρᾶν οὐ πρέπει τἀμήχανα.

I will stop only when I no longer have the strength.
Yet one should not begin by pursuing the impossible.

Sophocles, *Antigone*, v 91–92

Contents

Abstract **vii**

Contents **xi**

List of Figures **xiv**

List of Tables **xxi**

1 Introduction **1**

2 PMFC Harvesting: Background and Interface Design **11**
2.1 Related works 18
2.1.1 From intermittent polarization to intermittent harvesting 19
2.1.2 Intermittent harvesting architectures for embedded power delivery 23
2.1.3 Toward scalable and observable intermittent PMFC operation 27
2.1.4 Battery-free EIS capability 31
2.2 Intermittent loading design and implementation . . . 33
2.2.1 Timing and input multiplexer 35
2.2.2 Current and voltage measurement front end . 40
2.3 Summary 59

3 Experimental Evaluation of Intermittent PMFC Harvesting **61**
3.1 Experimental setup and measurement protocol . . . 63
3.1.1 PMFC construction and start-up 63

3.1.2 Loading protocols and power acquisition . . . 64
3.1.3 EIS measurement protocol 65
3.1.4 Measurement timeline and hydration events . 69
3.2 Power response to intermittent harvesting 69
3.2.1 Shunt-burden correction of PMFC-side power measurements 69
3.2.2 Preliminary timing exploration 71
3.2.3 Nine-day continuous and intermittent loading experiment 73
3.2.4 Discussion of power trends 76
3.3 Impedance response to loading history 77
3.3.1 Equivalent-circuit model for impedance interpretation 78
3.3.2 Effect of electrical loading 85
3.3.3 Effect of hydration 90
3.3.4 Impedance-analysis summary and outlook . . 95
3.4 Conclusions 98

4 CapDYN: Dynamic-Capacitance Architecture 101
4.1 Dynamic capacitance: related works and motivation 106
4.1.1 Differences from the State-of-the-Art 107
4.1.2 Centralized and federated solutions 117
4.1.3 Other relevant works 118
4.2 *CapDYN* hardware design 119
4.2.1 The progressive capacitance scaling logic . . 121
4.2.2 Main bank design: switching strategy and capacitor dimensioning solution 123
4.2.3 Power supply design and input/output power flow management 128
4.3 *CapDYN* software design 132
4.3.1 Execution Model and Definitions 133
4.3.2 Assumptions, Scope, and Limitations 135
4.3.3 *CapDYN* Task-based, Energy-Aware Runtime 137
4.4 Hardware and Software Implementation 138
4.4.1 Hardware Implementation 140
4.4.2 LTspice simulator 144
4.4.3 Software Implementation 145
4.4.4 Energy Profiling of Tasks and Task-to-Stage mapping 147
4.5 Summary 149

5 Experimental and System-Level Evaluation of CapDYN 151
5.1 Hardware evaluation 153
5.1.1 Efficiency measurement setup 154
5.1.2 Hardware validation of stage switching and power-flow control 157
5.1.3 Efficiency metrics computation 162
5.1.4 Infrastructural power, efficiencies, and losses as a function of current 165
5.1.5 Software and hardware controlled switching . 179
5.2 Application-level LTspice evaluation 180
5.2.1 Target battery-free sensing workload 181
5.2.2 Evaluation methodology 183
5.2.3 Experimental results 190
5.3 Discussion 200
5.4 Conclusions 206

6 Conclusions 209

A Appendix 215
A.1 Post-Switch Voltage and Ideal Switching-Loss Computation 215
A.1.1 General notation 215
A.1.2 D1 post-switch voltage 216
A.1.3 D2 post-switch voltage and load-range condition 216
A.1.4 Ideal charge-sharing loss 217

List of Figures

2.1 Operating principle of a PMFC, showing plant-assisted substrate supply, microbial oxidation at the anode, oxygen reduction at the cathode, ion transport through the soil, and electron flow through the external circuit. 12

2.2 Anode and cathode potential evolution of two PMFCs during start-up, measured with respect to an Ag/AgCl reference electrode. The 600 mV threshold indicates the open-circuit cell voltage above which the PMFCs were considered ready for power-production tests. 15

2.3 Intermittent harvesting architecture, featuring interleaved extraction from two PMFCs, current and voltage measurement front end. 33

2.4 PCB implementation of the intermittent-harvesting architecture. The bipolar supply is located in the top-left section of the board. The top-right section implements the current-sensing front end, whereas the bottom-right section implements the voltage-sensing front end. The timer element is located in the bottom-left section. 34

2.5 Schmitt trigger relaxation oscillator used as timer for the proposed intermittent harvesting interface. . . . 35

2.6 Current measurement front end schematic. 40

2.7 Feedback resistor scaling effect on output voltage spectral density of the current sense amplifier. 46

2.8 Shunt resistance effect on output voltage spectral density of the current sense amplifier. 47

2.9 Shunt capacitance effect on output voltage spectral density of the current sense amplifier. 48

2.10 Gain effect on output voltage spectral density of the current sense amplifier. 50

2.11 Validation of the developed current-sensing front end during intermittent harvesting from two PMFCs. The current estimated from the front end is compared with a reference measurement acquired using a Keithley DMM7510. 54

2.12 Power consumption of the reference BQ25505-based PMFC harvesting interface and of the proposed intermittent harvesting interface as a function of supply voltage. The proposed interface builds on the BQ25505 stage by adding the bipolar supply, intermittent timer, and optional voltage- and current-sensing front ends, which can be disabled to reduce power consumption when monitoring is not required. 58

3.1 Intermittent harvesting setup used for Cells I and H. 66

3.2 EIS measurement setup used in parallel with the intermittent harvesting setup shown in Fig. 3.1. 68

3.3 Effect of the shunt-burden correction on the PMFC-side electrical power of the intermittently loaded cells during the 9-day experiment. 70

3.4 Preliminary intermittent-harvesting frequency exploration over 6000 s, comparing 10 s, 100 s, and 1000 s periods at 50% duty cycle. 72

3.5 Extended preliminary frequency exploration over approximately 12 h, comparing 5 s, 10 s, and 20 s periods at 50% duty cycle. 73

3.6 Power evolution during the 9-day loading experiment. Cells I and H were operated intermittently, while Cell C was operated continuously. The intermittent period was 20 s during the first six days and 10 s during the last three days. 75

3.7 Effect of the shunt-burden correction on the PMFC-side electrical power of the intermittently loaded cells during the 9-day experiment. Dashed blue traces show the instrumented PMFC-side power, $P_{\mathrm{PMFC,inst}}$, obtained with the 6.7 Ω low-side shunt present. Solid red traces show the corrected estimate, $P_{\mathrm{PMFC,corr}}$, corresponding to the same equivalent external load after removing the shunt effect. The close overlap between pre- and post-correction traces indicates that the correction is small and does not change the observed power trends. 75

3.8 Temporal evolution of selected fitted anode parameters for the unloaded cell (U), continuously loaded cell (C), and intermittently loaded cells (H and I). The panels show the series resistance R_s, first-branch resistance R_1, first-CPE coefficient Q_1, first-CPE exponent n_1, and first-branch characteristic frequency f_{c1}. 85

3.9 Representative anode EIS fits showing the temporal evolution between Day 0 and Day 6 under continuous and intermittent loading. The continuously loaded anode shows a stronger temporal shift, whereas the intermittently loaded anode remains comparatively closer to its initial response. 87

3.10 Representative late-state anode EIS fits for unloaded, continuously loaded, and intermittently loaded cells. The comparison illustrates the effect of loading history on the anodic impedance response at Day 6. . . 89

3.11 Representative EIS fits showing the immediate effect of watering between Day 3 and Day 3W for the unloaded cell. The anode and cathode respond differently, highlighting the electrode-specific impact of hydration on the fitted impedance spectra. 92

3.12 Representative EIS fits showing the delayed post-watering response of Cell I between Day 1 and Day 2. The anode response shows changes in both the high-frequency/contact region and the low-frequency bounded-transport region. The cathode response shows a shift of the high-frequency intercept and first interfacial arc toward lower real impedance after watering. 93

4.1 The *progressive scaling* capacitor reconfiguration pattern for *CapDYN*, in two bank design variants. The initial configuration of both banks is *stage 1*, where all the capacitors are connected in series. As the stored energy increases, each design differently reconfigures the connections between the capacitors in successive steps to connect them in parallel. At *stage N*, all the capacitors of both banks are connected in parallel. Reproduced from [1]. 120

4.2 *CapDYN* block diagram (a), bank schematics (b) and input/output logic (c). Adapted from [1]. 127

4.3 Overview of the proposed energy-aware sensor architecture. The system integrates an energy harvester, a dynamically reconfigurable energy storage subsystem (*CapDYN*), and a computing unit interfacing with generic sensing and communication modules. The energy state interface enables coordination between the storage subsystem and the computing unit to support energy-admissible task execution. Adapted from [2]. © 2025 IEEE. 132

4.4 Energy-admissible task execution flow implemented through a hardware–software event-driven interaction between the scheduler (left) and CapDYN (right). The scheduler configures the required energy state for the next task and initiates execution only when CapDYN signals that the requested state has been reached. Task execution and power gating are coordinated via event signals. Reproduced from [2]. © 2025 IEEE. 139

4.5 Example application code using the *CapDYN* runtime. Adapted from [1]. 145

4.6 Code snippet that illustrates how system initialization and task transitions are managed within the *CapDYN* runtime. The `transition_to`() function sets up the HW register for the next task (line 14), checks if the required stage is achieved (line 15) and either executes the next task (line 16-20) or orders *CapDYN* to cut the load's supply (line 21-23). Adapted from [1]. 146

5.1 *CapDYN* efficiency measurement setup. On the top right, the source and sink SMUs, Keithley's 2470 and 2450, respectively. At the center, we can see National Instrument's NI-USB6259 data acquisition system. On the left, Agilent's 3631A dual power supply for the TL081P buffers (on the blue breadboard). 156
5.2 PCB of the CapDYN board, integrating the D1 bank design. 156
5.3 PCB of the D2 bank design. 157
5.4 *CapDYN* prototype under development. The central PCB integrates the entire control logic, as well as the D1 bank. The PCB on the left integrates the D2 bank. The protoboard on the bottom includes the TLV3691 comparators used to interface progressive scaling mechanism and bank analog switches. At the top, the TL081P buffers, while on the right we see the $\texttt{EN}_2$ actuated switch to trigger the input/output SMUs. 158
5.5 Charge/discharge voltage traces for the 90 μA reference test. 160
5.6 Power consumption of CapDYN with 3, 4, 5 stages, matched with either D1 or D2 bank. 166
5.7 Percentage of input energy lost to infrastructural consumption. 167
5.8 Cold start efficiency metrics of *CapDYN* as a function of charge/discharge current. At the top, the entire infrastructure's cold start efficiency is displayed, while at the bottom we focus only on the bank's cold start efficiency. 169
5.9 Percentage of input energy lost to leftover energy as a function of charge/discharge current. The top panel includes the full system, while the bottom panel focuses on the bank only. 170
5.10 Regimen efficiency of *CapDYN* as a function of charge/discharge current. At the top, the entire infrastructure's regimen efficiency is displayed, while at the bottom we focus only on the bank's regimen efficiency. . . . 173
5.11 Parasitic input energy losses versus charge/discharge current, normalized to total cold-start, bank cold-start, and bank regimen input energy. 174

5.12 Switching equalization transient losses versus charge/discharge current for the D1 bank, normalized to: total cold-start input energy, bank cold-start input energy, and bank regimen input energy. 178
5.13 Stage-energy availability of *CapDYN* as a function of capacitor-bank topology and total number of stages N. Solid lines and empty markers denote the D1 bank, while dotted lines and full markers denote the D2 bank; colors identify the relative stage position from stage N down to stage $N-6$. 196
5.14 Stage capacitance of *CapDYN* as a function of capacitor-bank topology and total number of stages N. Solid lines and empty markers denote the D1 bank, while dotted lines and full markers denote the D2 bank; colors identify the relative stage position from stage N down to stage $N-6$. 197
5.15 Voltage and task-execution comparison for the best-performing *CapDYN* configuration on the exponential task set. The figure compares the single-capacitor baselines S0, S1, and S2, the ablation baseline AB, and *CapDYN* with a four-stage D2 bank. This configuration provides the best overall tradeoff for the exponential task set, combining the highest effective efficiency and timely successful-task yield with earlier task completions relative to the single-capacitor baselines and to the ablation. 203
5.16 Voltage and task-execution comparison for the best-performing *CapDYN* configuration on the linear task set. The figure compares the single-capacitor baselines S0, S1, and S2, the ablation baseline AB, and *CapDYN* with a three-stage D2 bank. This configuration provides the best overall tradeoff for the linear task set, combining the highest effective efficiency and timely successful-task yield with earlier task completions relative to the single-capacitor baselines and to the ablation. 204

List of Tables

2.1 Summary of the current-sensing front-end noise analysis in the 0.1–5 Hz band. All measured ratios are normalized to the nominal front-end configuration: original feedback network, $R_S = 1\ \Omega$, no added shunt capacitance, and $G_I = 295$, under the same operating condition. 51

3.1 Active-sample statistics of the shunt-burden correction for the intermittently loaded cells during the 9-day experiment. 71
3.2 Average PMFC-side electrical power during the 9-day experiment before and after shunt-burden correction. The pre-correction values correspond to the instrumented PMFC-side power $P_{\mathrm{PMFC,inst}}$, whereas the post-correction values correspond to the corrected estimate $P_{\mathrm{PMFC,corr}}$. The intermittent period changed from 20 s to 10 s on Day 7, while Cell C remained continuously loaded. 74
3.3 Finite open Warburg model fitting, anode. 83
3.4 Finite open Warburg model fitting, cathode. 84

4.1 Comparison between *CapDYN* and representative prior work in storage architectures. Adapted and extended from [1]. 107

5.1 Experimental current range for each implemented configuration. 155
5.2 Stage-wise equivalent capacitance and available energy. 161
5.3 Description of the implemented tasks 183

5.4 Task power, energy, execution time, and matched *CapDyn* stage. Tasks and their respective functions are summarized in Table 5.3. ***OP*** stands for "over provisioning'. 184
5.5 Per stage capacity/energy vs number of available *CapDYN* stages. Single cap represent the baseline with fixed capacitance/energy. 184
5.6 Efficiency penalty of the unaware single-capacitor baseline S0 and effective-efficiency comparison with *CapDYN*. 191
5.7 Successful-task throughput and relative throughput gain of *CapDYN* with respect to the reference baselines. 192
5.8 Timely successful-task yield gain of *CapDYN* relative to the reference baselines. Positive values indicate higher time-weighted useful task yield. 193
5.9 Effective-efficiency comparison of *CapDYN* against the energy-aware baselines and ablation. 194
5.10 Aggregate first-completion advance of *CapDYN* relative to the reference baselines. Positive values indicate earlier first successful task completion under *CapDYN*. 195
5.11 Per-task first-completion shift of *CapDYN* relative to the ablation baseline AB. Negative values indicate earlier first task completion under *CapDYN*. 201
5.12 LTspice infrastructural power consumption of the single-capacitor baselines and of the reconfigurable *CapDYN* hardware used by both *CapDYN* and AB. The values are conservative with respect to the hardware measurements of Sec. 5.1.4, as they slightly overestimate infrastructural power while preserving the same topology- and stage-dependent scaling trends. 202

Chapter 1

Introduction

Battery-free embedded systems are a promising alternative to conventional battery-powered wireless sensor nodes, especially in applications where maintenance, replacement cost, and environmental impact are critical constraints. This is particularly relevant for precision agriculture and environmental monitoring, where sensing nodes may need to operate for long periods in remote or infrastructureless locations. In these scenarios, replacing or recharging batteries is often impractical, while deploying large numbers of battery-powered devices introduces sustainability concerns related to material use, chemical waste, and maintenance logistics. Energy-harvesting systems address these limitations by scavenging energy from the surrounding environment and using it to execute sensing, computation, and communication tasks without relying on a primary battery [3, 4, 5].

The transition from battery-powered to battery-free operation, however, changes the way embedded systems must be designed. A battery-powered node is typically engineered around a relatively stable energy reservoir, and the main design objective is to minimize average power consumption in order to maximize battery lifetime. A battery-free node, instead, must operate with an energy source that can be weak, variable, and strongly dependent on environmental conditions. The system stores harvested energy in a capacitor or supercapacitor and executes tasks only when enough energy is available. As a result, execution becomes intermittent: the node alternates between charging phases, active computation, and possible

power failures [6, 7]. In this operating regime, the relevant design question is not only how much energy can be harvested, but also when that energy becomes available, how efficiently it can be stored, and how effectively it can be converted into useful task completions.

This thesis addresses a central problem in battery-free embedded systems: energy harvesters often provide power that is low and variable, whereas embedded workloads require reliable bursts of energy to complete sensing, processing, communication, or monitoring tasks. This mismatch can be quantified by considering the time required to accumulate the energy needed by representative tasks. If the average harvested input power is in the range of 10–100 μW, which is representative of the PMFC power range characterized in Chapter 3, representative task energies translate into accumulation times ranging from seconds to hours: approximately 1–100 s for low-energy sensing tasks, 10–1000 s for wireless transmissions, and up to several hours for more demanding electrochemical measurements such as complete Electrochemical Impedance Spectroscopy (EIS) acquisitions.

Therefore, the main challenge is not only low average harvested power, but temporal energy buffering. The node must accumulate energy during idle intervals and release it in short bursts when sensing, communication, or impedance-measurement tasks are executed. This requirement creates two coupled design problems along the energy path. The first is a source-side extraction problem: the harvesting interface must extract energy from a weak and variable source, while possibly improving the average harvested power and reducing its variability. The second is a storage-side delivery problem: the energy buffer must make harvested energy available at the right time and support tasks with different energy and timing requirements. This thesis addresses these two problems through complementary contributions: controlled intermittent extraction from Plant Microbial Fuel Cells (PMFCs), and autonomous dynamic capacitance for adaptive storage and delivery. Their common goal is to increase the fraction of harvested energy that becomes completed battery-free tasks.

The source-side investigation focuses on Plant Microbial Fuel Cells (PMFCs) as sustainable bioelectrochemical energy sources for low-power environmental sensing. PMFCs exploit the interaction between plants, microorganisms, electrodes, and the surrounding soil to generate electrical power from naturally occurring biological

and electrochemical processes. Their appeal lies in their low material cost, potential long-term operation, and compatibility with agricultural environments. Unlike photovoltaic or radio-frequency energy harvesting, PMFCs can operate without direct illumination or external transmitters, making them particularly attractive for soil-embedded sensing nodes and smart-agriculture deployments. At the same time, PMFCs remain challenging power sources: their voltage is low, their power output is limited, their recovery dynamics are slow, and their electrical behavior depends on biological, electrochemical, and environmental conditions [8, 9].

These challenges motivate controllable extraction as a first requirement for PMFC-powered systems. PMFC energy extraction is not necessarily maximized by continuously loading the cell at a fixed operating point. PMFCs exhibit capacitive behavior, recovery dynamics, and loading-history dependence: excessive or poorly timed extraction can move the cell away from favorable operating conditions, while recovery intervals can allow voltage and concentration gradients to partially relax [10]. These dynamics also make conventional Maximum Power Point Tracking (MPPT) strategies designed for faster energy sources difficult to apply directly, suggesting that PMFC power management should consider not only instantaneous maximum-power operation, but also maximum-energy extraction over repeated harvesting and recovery cycles [9]. A promising direction is suggested by prior studies on *intermittent loading*, where loaded phases are alternated with unloaded recovery phases. These studies have shown that such operation can affect current generation, biofilm behavior, and mass-transport dynamics in microbial fuel cells and electroactive biofilms [11, 12, 13, 14, 10]. In PMFC energy harvesting, intermittent loading corresponds to alternating energy-extraction intervals, during which the cell is connected to the harvesting interface or load, with recovery intervals, during which it is left unloaded. This strategy may improve average harvested power by allowing partial recovery between extraction phases, while also reducing continuous electrical stress on the electroactive biofilm.

However, existing approaches leave a gap between electrochemical control and deployable harvesting. Three-electrode intermittent-polarization studies provide valuable insight into biofilm dynamics, charge storage, and electrode conditioning, but they do not directly implement usable energy harvesting [14, 15, 10]. Conversely, two-electrode intermittent-harvesting approaches are closer to deploy-

able PMFC-powered nodes, but often tie the harvesting and recovery intervals to capacitor charging, cell-voltage thresholds, or source recovery dynamics [11, 16, 17, 13, 18, 19]. As a result, the harvesting frequency and duty cycle are not controlled independently, but arise from the coupled dynamics of the PMFC, storage element, converter, and load. This makes it difficult to isolate how the imposed intermittency parameters affect PMFC power output and electrochemical behavior. This motivates a two-electrode harvesting interface that can impose intermittent extraction explicitly, while remaining compatible with converter-based energy delivery and with the stringent power budget required by PMFC-powered operation.

A second requirement for PMFC-powered systems is self-contained observability. Because PMFC behavior depends on soil, plants, hydration, microbial activity, and loading history, it must be monitored in situ to understand long-term operation effects. Battery-powered monitoring infrastructure would weaken the sustainability motivation by reintroducing maintenance requirements and environmental impact. This motivates self-powered or battery-free monitoring nodes that can observe PMFC behavior while drawing energy from the same class of sustainable sources they are designed to study [20].

The storage-side investigation focuses on dynamic capacitance as a way to adapt the energy buffer to intermittent battery-free execution. In conventional designs, the storage capacitance is fixed at design time, usually according to the most energy-demanding operation. This creates a fundamental tradeoff: a large capacitor can sustain longer tasks, but delays cold start and reduces responsiveness for shorter operations, while a small capacitor charges quickly but may not provide enough energy for demanding tasks. As a result, a fixed capacitor is poorly matched to heterogeneous workloads, where sensing, processing, communication, or electrochemical monitoring require different energy budgets and different response times.

Dynamic capacitance addresses this limitation by changing the equivalent storage capacitance during operation. The storage can start from a small effective capacitance to reach operating voltage quickly, and then progressively increase its capacitance as more energy becomes available. During discharge, the equivalent capacitance can be reduced again by reconfiguring the bank toward more

series-oriented stages. This produces a post-switch voltage boost, keeping the bank voltage above the load lower operating threshold for a longer interval and minimizing the residual energy that would otherwise remain unusable below the load voltage limit. In this way, dynamic capacitance does not only improve end-to-end efficiency; each capacitance step also represents an intermediate energy state that can be exposed to the runtime and matched to task requirements.

Existing approaches only partially address this problem. Software-controlled capacitor banks provide flexibility and can match storage size to task energy, but they require active control, state-of-charge monitoring, or microcontroller intervention, which increases overhead and couples reactivity to power consumption [21, 22]. Autonomous storage approaches reduce control overhead, but typically provide only a small number of capacitance levels or focus on collecting excess energy rather than matching task-specific energy requirements. Federated storage architectures distribute energy across peripherals, but can fragment energy across multiple capacitors and still require task-level management or lossy charge transfers [23, 24].

This motivates an autonomous, low-power dynamic-capacitance architecture that combines progressive capacitance scaling with runtime-visible energy states. This requirement is especially relevant for PMFC-powered systems, where the available power is limited: any overhead spent on storage management directly competes with the energy available for sensing, monitoring, and communication. The storage subsystem should therefore adapt its effective capacitance with minimal control cost, maintain a usable voltage during reconfiguration, reduce leftover energy after task execution, and expose enough information to the runtime to enable energy-admissible execution. In this model, tasks are released only when the current storage state guarantees that enough energy is available for completion.

In the remainder of this thesis, the term intermittency is used in two senses. Extraction intermittency refers to the deliberate alternation between energy-extraction intervals and recovery intervals imposed by the harvesting interface, as developed in Chapter 2 and evaluated in Chapter 3. Load intermittency refers to the discontinuous operation of the embedded node, which alternates between charging and task-execution phases because stored energy is not always sufficient for continuous operation, as addressed in Chapters 4

and 5.

The methodology adopted in these chapters follows the energy path of a battery-free node, from PMFC power extraction to storage-side task delivery. The objective is to understand how weak and variable harvested energy can be controlled at the source, buffered by adaptive storage, and converted into completed sensing, computation, communication, or monitoring tasks. The PMFC interface determines how energy enters the node, while the dynamic-capacitance subsystem and runtime determine when that energy becomes available for reliable task execution.

On the source side, Chapter 2 develops a two-electrode intermittent harvesting interface for PMFCs. The key methodological choice is to impose extraction and recovery intervals explicitly, instead of allowing them to emerge from capacitor thresholds or load demand. This makes period and duty cycle controlled circuit-level variables while preserving compatibility with converter-based energy delivery. The interface is characterized in terms of timing operation, voltage and current measurement, and power consumption. In Chapter 3, the interface is then used to compare continuous and intermittent extraction from PMFCs, while Electrochemical Impedance Spectroscopy (EIS) is used as a supporting observation tool to assess whether different loading histories produce measurable electrochemical signatures.

On the storage side, Chapter 4 develops *CapDYN*, an autonomous dynamic-capacitance architecture for battery-free intermittent systems. Instead of relying on a single fixed storage capacitor, *CapDYN* reconfigures a capacitor bank during charge and discharge, changing the equivalent capacitance seen by the harvester and by the load. This creates discrete storage stages that correspond to different usable energy levels: smaller effective capacitances improve recharge time and reactivity for low-energy tasks, while larger effective capacitances provide the energy required by more demanding operations. Chapter 5 first evaluates the hardware layer in terms of capacitor-bank reconfiguration, leftover energy, conversion losses, and infrastructural power. It then assesses the system-level impact through an energy-aware runtime that maps profiled tasks to storage stages and dispatches each task only when the current stage contains enough energy to complete it. In this way, *CapDYN* is assessed both as an autonomous storage circuit and as a hardware–software mechanism for reducing wasted energy, avoiding failed executions,

and making harvested energy available to applications earlier.

The technical connection between the two contributions is particularly relevant for PMFCs because their low, slowly varying, and history-dependent power must ultimately supply embedded workloads with heterogeneous and bursty power and energy requirements. *CapDYN* addresses this mismatch by adapting the storage state to task energy demand, allowing lower-energy tasks to be executed with shorter charging delays while releasing more demanding tasks only when sufficient energy has been accumulated. In an integrated implementation, the intermittent-harvesting strategy would determine both the average power and the temporal profile delivered by the PMFC: higher average harvested power would shorten *CapDYN*'s charging times, making task-matched storage states available more frequently and thereby increasing the throughput of tasks dispatched by *CapDYN*'s runtime.

The results reported in the following chapters support this source-to-load perspective through separate source-side and storage-side evaluations. On the source side, reducing the intermittent period from 20 s to 10 s coincided with average-power increases of 159% and 248% for the intermittent PMFCs, compared with 31% for the continuous reference. Relative to the continuous cell, Cell I improved from +44% in the initial 20 s phase to +185% in the final 10 s phase, while Cell H shifted from −41% to +57%. These gains support controlled intermittent extraction at properly selected periods as a promising strategy, while the simultaneous increase of the continuous reference indicates that biological and environmental dynamics remain important confounding factors. The timer-based intermittent extension adds 6.8–7.9 μW with respect to the BQ25505-only reference interface, or 6.5–6.8 μW with respect to the BQ25505-plus-bipolar baseline; the sensing front ends add further overhead but can be disabled outside characterization and power tracking.

On the storage side, *CapDYN* validates autonomous charge/discharge reconfiguration with measured control-infrastructure consumption below 10 μW, and reduces leftover energy from 23.67% of the input energy in the fixed-capacitor baseline to 0.16%–5.57% across the measured configurations. In the LTspice-based application-level evaluation, the energy-aware runtime exploits the intermediate energy states exposed by the hardware to avoid failed executions and make useful task completions available earlier, improving effective efficiency by up to 34.10% and achieving aggregate first-

completion advances of up to 45.00% with respect to the evaluated single-capacitor baselines. The ablation benchmark further shows that these temporal gains come from coupling the reconfigurable bank with energy-admissible dispatch: the bank exposes intermediate energy states, while the runtime turns them into earlier task releases. Together, these results show how controlled extraction from PMFCs and adaptive storage address the same system-level objective: increasing the fraction of scarce sustainable energy that becomes completed battery-free tasks.

The remainder of the thesis is organized around two complementary parts of the battery-free energy path, each separated into a design-oriented chapter and an evaluation chapter.

Chapter 2 focuses on PMFCs as sustainable harvesting sources and on the design of a controlled intermittent power-extraction interface. The chapter first introduces the operating principle of PMFCs and discusses their main limitations as practical power sources for battery-free embedded systems: low voltage, low power, slow recovery dynamics, biological variability, and limited direct observability. It then reviews related work on intermittent microbial fuel cell operation in Sec. 2.1, distinguishing between electrochemical studies based on three-electrode control and embedded harvesting approaches based on capacitor or converter dynamics. This motivates the need for a two-electrode interface that can impose extraction and recovery timing explicitly while remaining compatible with energy-harvesting converters. The chapter then presents the design and implementation of the intermittent loading platform in Sec. 2.2. This includes the timing and input-multiplexing stage, the voltage- and current-sensing front ends, the noise and calibration analysis of the current-sensing chain, and the power-consumption characterization of the complete interface.

Chapter 3 experimentally evaluates this platform using real PMFCs. The experimental setup and measurement protocol are described in Sec. 3.1; the power response to intermittent harvesting is analyzed in Sec. 3.2, including the preliminary timing exploration and the nine-day continuous and intermittent loading experiment. Finally, Sec. 3.3 uses EIS measurements to investigate whether different loading histories produce different anode and cathode impedance responses, opening the possibility of future harvesting strategies that combine power-based and electrochemical feedback.

Chapter 4 moves from source-side extraction to storage and load delivery. It presents *CapDYN*, an autonomous dynamic-capacitance architecture for battery-free intermittent systems. The chapter first motivates the need for dynamic storage, explaining why fixed capacitors are poorly matched to heterogeneous workloads and why existing reconfigurable or federated storage approaches still leave an efficiency–reactivity tradeoff. The comparison with prior work is developed in Sec. 4.1, leading to the design requirements for autonomous dynamic capacitance: fine-grained energy states, low leftover energy, efficient capacitance transitions, low infrastructural consumption, and immediate reactivity to harvested-power dynamics.

The hardware design of *CapDYN* is presented in Sec. 4.2. It describes the progressive capacitance-scaling logic in Sec. 4.2.1, the capacitor-bank topologies and switching strategy in Sec. 4.2.2, and the power-supply and input/output flow management in Sec. 4.2.3. The software layer is introduced in Sec. 4.3: Sec. 4.3.1 defines the execution model and the energy-admissibility condition, while Sec. 4.3.3 presents the runtime that maps profiled tasks to storage stages and executes them only when the required energy state is available. Implementation details are reported in Sec. 4.4, including the hardware prototype, LTspice simulator, software implementation, and task-energy profiling.

Chapter 5 evaluates *CapDYN* at circuit and application level. It first characterizes the storage hardware, measuring switching behavior, efficiency, infrastructural power, and loss mechanisms, with the main efficiency and power results reported in Sec. 5.1.4. The analysis then moves to application-level operation in Sec. 5.2, using an LTspice digital twin, real task-energy profiles, and harvested-power traces to quantify how energy-state-aware execution affects effective efficiency, responsiveness, and task completion. Finally, conclusions are presented in Chapter 6.

Overall, Chapters 2 and 3 address how energy can be extracted from and observed in a sustainable but variable source, whereas Chapters 4 and 5 address how that energy can be stored and converted into useful task execution. Their common contribution is a source-to-load view of battery-free design, in which harvesting, storage, monitoring, and scheduling are treated as parts of the same energy-management problem.

Chapter 2

PMFC Harvesting: Background and Interface Design

Parts of this chapter are adapted from the author's previously published works on Plant Microbial Fuel Cells and battery-free electrochemical monitoring [20, 9, 25, 26].

Precision farming increasingly relies on wireless sensor networks (WSNs) for continuous environmental monitoring, enabling data-driven control of crops, soil conditions, and key resources such as water and pesticides. In this context, battery-free sensing nodes are attractive because they can reduce maintenance requirements and environmental impact by operating from harvested energy rather than from replaceable batteries [3, 27, 28]. Thanks to advances in ultra-low-power hardware and intermittent-computing paradigms [29, 30, 31, 7], these systems can scavenge energy from the environment, store it in a capacitor or supercapacitor, and execute sensing or communication tasks when sufficient energy is available.

Within this design space, Plant Microbial Fuel Cells (PMFCs) offer a particularly attractive path toward sustainable agricultural sensing. Unlike photovoltaic, wind, or thermoelectric harvesters [32, 33, 34], PMFCs exploit the plant–soil–microorganism system itself as an energy source, making them well matched to soil-embedded and long-term monitoring applications. However, their output power

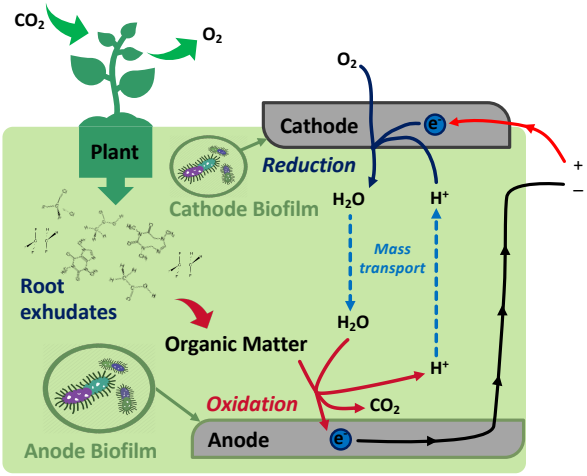


Figure 2.1: Operating principle of a PMFC, showing plant-assisted substrate supply, microbial oxidation at the anode, oxygen reduction at the cathode, ion transport through the soil, and electron flow through the external circuit.

is lower and more variable than that of established ambient harvesters [35]. Their voltage is typically sub-volt, their recovery dynamics are slow, and their behavior depends on biological, electrochemical, and environmental conditions. For this reason, the main question addressed in this chapter is how PMFC energy can be extracted and observed in a controlled way while remaining compatible with deployable battery-free electronics.

In its simplest form, shown in Fig. 2.1, a PMFC consists of an anode and a cathode embedded in soil at different depths, with a suitable plant growing above them. The anode is typically placed in the anaerobic region close to the rhizosphere, whereas the cathode is located in a more oxygenated region, usually closer to the soil surface. PMFC operation relies on the metabolic activity of naturally occurring exoelectrogenic bacteria and on their symbiotic interaction with the plant, which continuously supplies organic compounds to the rhizosphere through photosynthesis-derived root exudates [36]. As microorganisms oxidize the organic matter available near the anode, they release electrons and ionic species, including protons. The electrons are collected by the anode and flow through the ex-

ternal circuit toward the cathode, where they can be used by the load or harvested by the power-management interface. In parallel, ionic species migrate through the soil, closing the internal electrochemical path. At the cathode, the reduction reaction completes the process. This reaction is typically oxygen reduction, although it can also be catalyzed by a biocathode [37], where aerobic electrotrophic microorganisms form a catalytic biofilm on the electrode surface, or involve alternative electron acceptors such as ferricyanide, nitrate, sulfate, or metal species [38]. Accordingly, the electrode materials are chosen to favor the development of electroactive microbial biofilms, which retain the active microorganisms and act as biocatalysts for the redox reactions occurring at both electrodes [39]. As illustrated in Fig. 2.1, the overall PMFC output results from the coupled action of plant physiology, microbial metabolism, electrode reactions, soil transport phenomena, charge transport through the soil, and the external electrical load.

PMFCs enable continuous low-power generation for as long as the plant remains alive [26]. Thanks to their simple architecture, PMFCs can be realized at very low cost using basic materials such as carbon-based electrodes, conductive wires, soil, and a plant, while still delivering power in the mW range [40], sufficient to supply suitably designed electronic loads employing aggressive power-management techniques [41]. Moreover, PMFCs have been reported to achieve power densities up to 18 times higher than conventional MFCs [42], mainly because the plant continuously enriches the soil with organic substrates, enabling long-term operation with minimal maintenance. Accordingly, PMFC performance is strongly influenced by plant physiological conditions, especially photosynthetic activity and rhizodeposition, which regulate substrate availability for electroactive microorganisms and ultimately affect the generated power [8].

Problem statement: unfortunately, despite their strong sustainability appeal, PMFCs still face several dependability issues that limit their widespread adoption as practical energy sources. Unlike conventional harvesters, PMFCs are living energy sources whose electrical behavior evolves together with the biological and electrochemical state of the system. Their performance results from the coupled action of plant physiology, microbial activity, electrode processes, and soil transport phenomena, whose interaction is still only partially understood. We group these challenges into three main

points.

① *Low voltage, low power, low energy*: a PMFC cell typically outputs sub-volt voltages, which are insufficient to directly supply most low-power electronics. Boosting circuits are therefore required to reach operative voltage levels, typically in the range from 1.8 V to 3.6 V, while power scale-up strategies still require further investigation. In practice, increasing PMFC output is not straightforward: size scale-up exhibits an optimum reactor size in terms of power density [43], whereas multi-cell coupling in series or parallel remains challenging because differences in output voltage, recovery dynamics, and internal impedance can create imbalances that hinder effective aggregation [44]. In addition, the highly capacitive and slow-varying nature of PMFCs makes maximum-power-point operation more difficult than in faster energy sources, calling for ad-hoc harvesting strategies [9]. In this context, intermittent operation has emerged as a promising direction, since periodically alternating harvesting and recovery phases may improve both the average harvested energy and long-term biofilm conditioning [11, 18, 19, 12, 17]. However, the implementation of intermittent power extraction in practical two-electrode energy-harvesting systems is still not fully understood, especially in relation to the more established three-electrode electrochemical control literature [14, 15, 10]. Indeed, most current insights on intermittency derive from three-electrode electrochemical control, whereas deployable PMFC harvesters operate as two-electrode systems, making the transfer of such knowledge nontrivial. Finally, the total energy budget extractable over the PMFC lifetime is constrained by *aging phenomena* such as biofouling and electrode inactivation.

② *Low stability*: once a PMFC is built, a highly variable *start-up* period, ranging from days to months, may elapse before appreciable power production begins (Fig. 2.2). Even after start-up, the PMFC output remains strongly dependent on plant physiological state, microbial activity, substrate availability, and electrode condition. Consequently, power production may fluctuate over time or progressively degrade because of aging phenomena such as biofouling, salt deposition, cathode alkalinization, anode acidification, and dead-layer formation on the electrodes. More generally, unpredictable behavior arises, including sudden biofilm detachment and even *voltage reversal* (VR), that is, the inversion of cell polarity [45]. At the same time, recent studies on intermittent loading suggest that

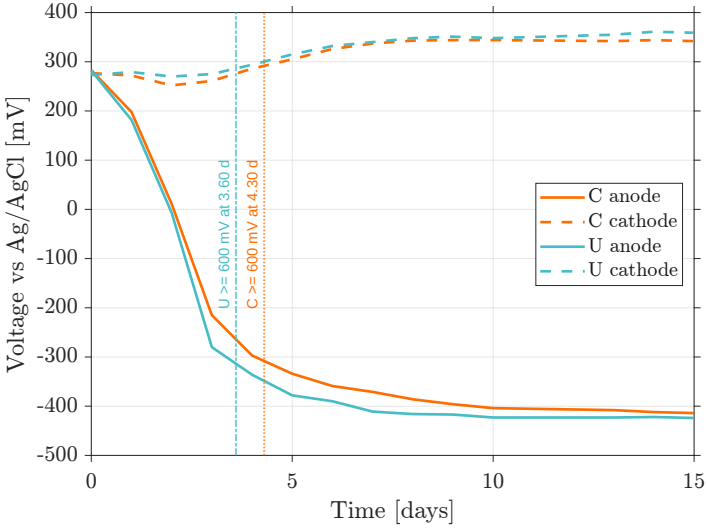


Figure 2.2: Anode and cathode potential evolution of two PMFCs during start-up, measured with respect to an Ag/AgCl reference electrode. The 600 mV threshold indicates the open-circuit cell voltage above which the PMFCs were considered ready for power-production tests.

the temporal profile of the electrical load can itself influence biofilm activity, mass transport, charge storage, and long-term power generation [14, 15, 46, 10]. PMFC instability is thus a manifestation of the evolving biological and electrochemical state of the system, which is also shaped by the adopted harvesting strategy.

③ *Destructive direct state observation*: the health of the bacteria, and the state of the electrodes and their surrounding substrate cannot be assessed *directly* without destructive sampling that significantly alters the PMFC state. Yet, these parameters critically determine future power output and long-term reliability. Indirect observation through electrochemical measurements is possible, for the sake of both biosensing and Maximum Power Point Tracking (MPPT), but it is in itself challenging. On the one hand, the energy source exhibits long transients, capacitive behavior, and loading-history dependence, which make the continuously evolving Maximum Power Point (MPP) difficult to reliably track. Conventional MPPT strategies, such as the FOCV methods commonly integrated in converters for faster sources, are thus poorly suited to PMFCs [9]. At the same time, electrochemical characterization must be compatible with the self-powered, battery-free constraints of the PMFC node. This is non-trivial because electrochemical characterization techniques require controlled excitation, signal acquisition, and processing, all of which introduce an energy overhead that is difficult to sustain in ultra-low-power systems. The challenge is even more pronounced in PMFC-powered nodes, where the energy source is weak, variable, and directly coupled to the electrochemical processes under observation. A practical diagnostic approach must therefore provide useful information about the PMFC state while keeping the measurement energy and hardware overhead low enough for battery-free operation.

Taken together, these three aspects point to two complementary needs: a harvesting strategy capable of improving PMFC power delivery while remaining compatible with battery-free energy extraction, and an indirect, non-destructive observation method capable of tracking the PMFC state within the same resource-constrained setting.

Chapter contributions: This chapter addresses these needs by considering two complementary capabilities for battery-free PMFC nodes: intermittent two-electrode operation and Electrochemical Impedance Spectroscopy (EIS)-based state observation. EIS is a

non-destructive electrochemical technique that applies a small sinusoidal perturbation over a range of frequencies to estimate the impedance of the cell. The resulting spectrum provides diagnostic information about charge transfer, mass transport, microbial activity, and reactor conditions [47, 48]. In previous work [20], we demonstrated that this measurement principle can be embedded into a PMFC-powered battery-free node, using the same electrodes for energy extraction and impedance sensing. That implementation validated an ultra-low-power EIS front end against a reference potentiostat and showed that a 19-point sweep from 21.3 mHz to 21.8 kHz requires 3.64 J, making impedance-based state observation compatible with PMFC-level energy production.

Within this framework, battery-free EIS provides the state-observation capability, while the complementary power-extraction capability developed in this chapter is a controlled two-electrode intermittent-harvesting interface. This interface is experimentally evaluated in Chapter 3, where impedance measurements are used to assess whether different loading profiles leave measurable electrochemical signatures in the PMFC response and whether such signatures could support future PMFC-specific energy-management strategies, including adaptive intermittent harvesting.

The main design contribution of this chapter is a two-electrode intermittent-harvesting interface that connects PMFCs to an energy-harvesting DC–DC converter while explicitly controlling the harvesting frequency and duty cycle. In contrast to capacitor-based or hysteretic approaches, where the intermittent timing is tied to the instantaneous recovery dynamics of the cell, the proposed interface imposes the timing externally. This enables intermittent harvesting to be studied as a controlled electrical treatment while preserving compatibility with embedded power delivery. In the implemented prototype, the timing and monitoring functions are realized with ultra-low-power components, allowing the interface to target battery-free operation.

The remainder of the chapter is organized as follows. Sec. 2.1 reviews the literature on intermittent MFC and PMFC operation, distinguishing three-electrode intermittent polarization from deployable two-electrode intermittent harvesting, and discusses existing approaches for scalable and observable PMFC energy management. Sec. 2.1.4 summarizes the previously demonstrated battery-free EIS capability that provides the system-level basis for embedded PMFC

state observation. Sec. 2.2 presents the proposed intermittent-harvesting interface, including the timing circuit, input-multiplexing strategy, and voltage/current monitoring front end. The experimental evaluation of the developed interface and of the resulting PMFC response is presented in Chapter 3.

2.1 Related works

Intermittent operation has emerged as a promising strategy to improve the power output of both MFCs and PMFCs, a particularly meaningful objective in the context of this thesis, where these systems are studied for the powering of embedded electronics. Unlike conventional energy harvesters, PMFCs are living electrochemical systems, and their output is not only determined by the instantaneous electrical load, but also by the state of the electroactive biofilm, the availability of substrate, the plant physiological conditions, and the transport phenomena taking place in the surrounding medium [8]. As a consequence, continuously loading the cell at a fixed operating point does not necessarily maximize either the instantaneous extracted power or the total energy recovered over time [49, 50, 12]. This has motivated a growing body of literature on *intermittent operation*, where harvesting phases are alternated with open-circuit phases.

A first important distinction must be drawn between *intermittent polarization* and *intermittent harvesting*. Intermittent polarization is typically implemented in a *three-electrode* configuration, in which the potential of a selected electrode is controlled with respect to a reference electrode (e.g., Ag/AgCl or Saturated Calomel Electrode) using a potentiostat/galvanostat. This configuration allows the electrochemical conditions experienced by the electrode and its electroactive biofilm to be prescribed directly, and is therefore particularly suitable for investigating the effects of electrode potential, polarization time, and recovery time on microbial metabolism, charge storage, and biofilm development. However, such potentiostatic control is primarily an electrochemical characterization and conditioning method and does not directly address the delivery of harvested energy to an embedded load [51].

Intermittent harvesting, instead, operates the MFC or PMFC as a *two-electrode* energy source and alternates energy-extraction

phases with recovery phases, for example by switching an external load [12], charging and discharging a storage capacitor within prescribed voltage windows [17, 11, 18], or modulating the input of a DC–DC converter [50]. In this configuration, the controlled electrical quantity is the terminal cell voltage rather than the potential of an individual electrode with respect to a reference. Consequently, the anode and cathode potentials are not independently imposed, but emerge from the full-cell operating condition and the combined electrode losses. The evolution of the electrode potentials is affected by electrode kinetics [52], substrate availability and mass transport [53], and local pH [54], among other factors.

This distinction is important when transferring results from the established three-electrode literature to practical intermittent harvesting. Even when the same nominal polarization/recovery times or duty cycle are applied, a two-electrode harvester does not necessarily reproduce the electrode-potential trajectories experienced under three-electrode control. Therefore, results obtained from intermittent polarization studies should be interpreted as providing mechanistic hypotheses and design guidance, rather than directly transferable optimal harvesting parameters. Three-electrode intermittent polarization studies have shown that the imposed polarization pattern can affect transient charge release, charge-storage behaviour, electroactivity, and longer-term biofilm structure and adhesion [14, 55, 10]. A two-electrode harvesting system must then determine whether these effects remain beneficial when the MFC is operated as an actual energy source.

This creates a central conceptual and technological gap. Most mechanistic understanding of intermittent operation originates from controlled three-electrode electrochemical studies, whereas deployable battery-free PMFC systems necessarily operate through two-electrode energy extraction. Bridging these domains therefore requires a harvesting architecture that retains explicit control over the temporal loading pattern while allowing the resulting electrode-level electrochemical response to be observed.

2.1.1 From intermittent polarization to intermittent harvesting

The main reason why intermittent operation may outperform continuous harvesting lies in the non-trivial dynamic response of the

electroactive biofilm, which has been investigated primarily through *intermittent polarization* studies. Although these experiments do not reproduce the coupled electrode conditions of a two-electrode harvester, they provide insight into the electrochemical and biological processes that may be exploited by alternating energy-extraction and recovery phases. During open-circuit, microbial metabolism can continue even though electrons are no longer drained through the external circuit: Ter Heijne *et al.* [15] examined the storage mechanisms of electrons during open-circuit conditions, showing that microorganisms can continue metabolizing even when electrons cannot be transferred to the external circuit. During open circuit, microbes can store electrons either in intermediate reduced cell components for short-term storage, such as cytochromes and flavins, or in polymeric substances for long-term storage. When the electrical circuit is closed again, the release of this stored charge gives rise to a transient current or power peak superimposed on the steady-state current associated with ongoing microbial metabolism [56]. This transient contribution is one of the main mechanisms behind the higher average harvested energy observed under intermittent loading [14, 10]. From this perspective, the PMFC behaves as a living electrochemical system with memory, whose response also depends on previous loading conditions and on the timescale of the imposed electrical perturbation. Guo *et al.* [14] discussed the tradeoff between allowing enough open-circuit time to accumulate stored charge and avoiding excessive reduction in microbial activity: in fact, intermittent polarization of *G. sulfurreducens* biofilms improves current output only within a limited timing window. Pulse widths shorter than 5 min increased current with respect to the 18-day baseline, with a maximum response of 1.9 mA at 0.5 min, whereas longer pulse widths reduced the benefit granted by intermittent operation. Guo's results therefore indicate that intermittent operation is strongly timing-dependent and involves a tradeoff between enhanced electron-transfer rate and reduced coulombic efficiency. Similarly, Zhang et al. [10] show that duty cycle directly tunes long-term EAB performance. Literature had often suggested optimal duty cycles around 75–95% [13], but this long-term experiment found that 50% duty cycle with 10 s polarization / 10 s open circuit produced the highest cumulative charge and best long-term EAB performance, at the cost of a slower initial biofilm growth. Taken together, these studies show that the temporal polarization pattern

is itself an electrochemical control variable: polarization time and recovery time (thus, duty cycle and frequency) can modify both the immediate electrode response and the longer-term behaviour of the electroactive biofilm. This provides a key motivation for controlled intermittent harvesting, in which these temporal variables should be explicitly imposed rather than allowed to emerge indirectly from capacitor charging or source-recovery dynamics. At the same time, the optimal timing values identified under three-electrode control cannot be transferred directly to a two-electrode harvester, because the individual electrode potentials are not independently imposed and the optimization objective may differ. The relevant question for harvesting is therefore whether comparable timing-dependent benefits persist when intermittency is applied while extracting usable energy from the complete PMFC.

Caizán-Juanarena *et al.* [57] complement the bioelectrochemical interpretation of intermittent power enhancement by drawing an analogy between MFC transient behavior and electrochemical capacitors. From this viewpoint, the biofilm-electrode interface acts as an internal charge-storage element: during open-circuit phases it can partially self-recharge, whereas upon reconnection it releases the accumulated charge as a transient power contribution. This interpretation is particularly relevant for PMFCs, whose electrodes and biofilms often exhibit strong capacitive behavior and long transient times. As a result, PMFCs should not be treated as fast sources suited to conventional maximum-power-point tracking alone [58]; rather, their optimization can be framed as either a maximum-power-point-tracking or a maximum-energy-extraction problem [9], where duty cycle, recovery time, and operative voltage become central design variables.

Another reason why intermittency is of great interest is that its effect may extend beyond short-term charge release and directly influence the long-term state of the biofilm itself. Several studies suggest that the temporal profile of the applied load can modify biofilm activity, spatial structure, mass transport, and even community composition [59, 46, 10]. For instance, intermittent operation has been associated with enhanced mass transport near the electrode, increased utilization of deeper biofilm layers, and changes in the preferred charge-storage pathways of electroactive microorganisms [46, 10]. Zhang *et al.* [10] further showed that periodic polarization affects not only short-term energy extraction, but also

anodic biofilm thickness, detachment behavior, and the properties of the electroactive bacterial community.

These observations introduce an important conceptual shift: intermittent operation should not be regarded merely as a circuit-level strategy for extracting more energy from an otherwise unchanged source. Rather, the imposed temporal loading profile can act as an electrochemical perturbation capable of influencing the biological and electrochemical state of an electroactive biofilm over time. For practical PMFC harvesting, this raises the possibility that the harvesting strategy may affect not only short-term power extraction but also longer-term source behaviour, potentially improving biofilm stability over time. Establishing whether the effects reported under controlled electrochemical conditions also occur under two-electrode PMFC harvesting therefore requires electrode-level observation during or following prolonged harvesting experiments.

Lastly, intermittent operation can also be leveraged for simultaneous *biosensing*. Fernandez-Gatell *et al.* [60] showed that the reconnection peak can last from the order of minutes up to tens of minutes and can be correlated with biomass and organic-matter content, COD concentration, and microbial activity, assessed through ATP content and enzymatic activity, within an MFC. Importantly, they highlighted that an open-circuit period of at least 5 min is required for the biosensing effect to become observable. Chung *et al.* [61] also correlated the reconnection current peak with the concentration of naphthenic acids in aqueous solutions. Intermittency can therefore be leveraged not only for energy extraction, but also as a repeated perturbation-response mechanism for sensing, since each harvesting cycle generates a transient current peak carrying information on the biological and chemical state of the system. This reinforces the interest of controlled intermittent operation: by explicitly setting recovery time, harvesting time, duty cycle, and operating point, the same architecture may support both improved power management, biofilm stabilization and non-destructive state observation, provided that the timing is chosen to satisfy the requirements of both energy harvesting and sensing.

2.1.2 Intermittent harvesting architectures for embedded power delivery

The studies discussed in the previous section mainly address *intermittent polarization*, typically in three-electrode setups that enable precise electrochemical control but do not allow direct energy harvesting. In this thesis, we aim to translate the knowledge gained from such controlled configurations to *intermittent harvesting*, where the PMFC operates as a two-electrode source and the extracted energy can be used to power embedded electronics. From an implementation standpoint, intermittent-harvesting strategies can be grouped into a few main families according to the state variable used to regulate the frequency and duty cycle of intermittent operation.

A first group comprises *capacitor-based approaches*, in which the MFC/PMFC charges an intermediate storage capacitor that is subsequently discharged into a load. In these systems, the capacitor naturally sweeps the I-V curve of the source during charging, while the source is disconnected during the capacitor discharge phase. Dewan *et al.* [11, 16] proposed a methodology for selecting the storage-capacitor size and charging thresholds, showing that both parameters exhibit an optimum in terms of average harvested power. In the same line of work, Beyenal *et al.* [16] showed that intermittent energy harvesting can yield up to 111% higher average power than continuous harvesting, thereby highlighting the potential of duty-cycled PMFC operation for powering embedded devices. Hsu *et al.* [17] implemented capacitor-based intermittent harvesting by cycling the storage capacitor between approximately 0.5 and 0.4 V, close to the MFC maximum-power region, and compared intermittent harvesting with continuous operation near 0.475 V. Although capacitor-driven schemes are practical, the resulting harvesting and recovery durations are determined by the capacitor value, the voltage window, and the instantaneous MFC power output. As a result, the control is naturally suited to embedded power delivery, but the effective timing of the harvesting cycle still depends on the PMFC dynamics, while the open-circuit interval is constrained by the discharge time of the storage capacitor, which typically corresponds to the execution time of the embedded task. Such a discharge interval may therefore remain too short for the PMFC to fully exploit the recovery mechanisms underlying intermittent-power enhance-

ment [14, 10]. Consistently with this methodological limitation, Hsu *et al.* [17] observed that the main benefit of intermittent harvesting lies in startup behavior, whereas optimized intermittent and continuous strategies tend to converge to similar steady-state power. Taken together, these observations indicate that the lack of steady-state improvement is likely not an intrinsic limitation of intermittent harvesting, but rather a consequence of recovery intervals that are too short to fully benefit PMFC operation over time.

A second group leverages *cell-voltage hysteretic control*, where the PMFC voltage itself is used as the control variable, rather than the capacitor voltage. In this case, the source is disconnected once a lower cell-voltage threshold is reached and reconnected only when the cell recovers above an upper threshold [13]. Wang *et al.* [18] implemented intermittent harvesting using a hysteretic controller that disconnects the MFC between two voltage thresholds, whereas Do Park *et al.* [19] introduced threshold adaptation through digitally controlled potentiometers driven by a DSP-based MPPT loop. Compared with capacitor-voltage control, this approach is more directly tied to the instantaneous electrochemical state of the PMFC, since the harvesting frequency and duty cycle adapt to the cell recovery dynamics themselves. This makes cell-voltage control attractive, as the intermittency timing can follow the state of the PMFC. At the same time, however, it becomes more difficult to isolate cause-and-effect relationships between duty cycle, frequency, and PMFC performance, because these timing variables are not independently imposed, but rather emerge from the evolving state of the cell.

A third group controls the cell operating point directly, through either *external loads or converters*. Liang *et al.* [62] investigated transient-state regulation by periodically connecting and disconnecting an MFC to an external resistor, explicitly studying the combined effect of frequency, duty cycle, and external resistance. They reported that power density increased with frequency, reached a maximum at 1 Hz, and peaked at a 75% duty cycle with a 200 Ω external resistance, achieving 1238 mW m^{-2}, 18.4% higher than steady-state operation. Coronado *et al.* [63] extended this approach using pulse-width-modulated external resistance over 0.1–1000 Hz, showing that high-frequency switching above 100 Hz increased MFC power output by 22–43% relative to constant-resistance operation; in longer R-PWM tests, a 95% duty cycle produced 3.6–3.8 mW, corresponding to 72–76 mW L^{-1}. Although these studies come close

to jointly optimizing timing and operating condition, the operating point is still imposed through an external resistor. While this is a simple and effective way to study transient MFC operation, practical PMFC power-management systems usually impose the cell voltage through a boost converter. Gardel *et al.* [12] move in this direction and show that duty cycling a 0.5 V programmable load in a 15-anode environmental MFC has an optimal timing window. Sequentially switching among anodes yielded the highest charge recovery at 15 s ON per anode, corresponding to approximately 210 s of OFF time and about 17 C/day. Shorter OFF intervals limited anode voltage recovery: below 10 s, recovery remained below 20% and current peaks halved. The authors attributed this behavior to diffusion-limited replenishment of electron donors in the biofilm and surrounding diffusion layer. While programmable fixed-voltage loads can emulate the operating point imposed by an input-modulated converter, and provide a useful tool for the study of intermittent operation on MFCs, they do not address the practical integration of ultra-low-power, intermittent energy conversion in an autonomous node. Literature remains comparatively sparse in this area. Some works implement converter-based intermittent extraction only indirectly, for example in the context of interleaving energy extraction from different units based on power ranking [50], without focusing on the effect of intermittency itself. Another work by Wang et al. [64] focuses on converter-based MFC harvesting and optimizes electrical parameters such as inductance, switching frequency, and converter duty ratio for power and energy production. However, the analyzed frequencies, which are in the kHz range, are relevant to converter switching operation rather than to the biological-scale frequencies of intermittent harvesting, which are most likely in the seconds/minutes range. As a result, the study addresses power-electronics optimization, but does not investigate the ON/OFF frequencies and recovery intervals required to exploit intermittency-induced improvements in MFC performance. More recent work by Fernandez Feito *et al.* [58] represents the state of the art in converter-based MFC power management, implementing a PMS with individual synchronous boost converters and MPPT incremental conductance control for up to five MFCs, and efficiencies up to 87%. This work is highly relevant because it addresses the practical integration of MFC energy into a usable power-management architecture and explicitly accounts for the slow dynamics of real tubular MFCs by adopting

a 60 s MPPT update interval. However, the controlled duty cycle is the high-frequency PWM duty cycle of the boost converters, rather than the biological-scale ON/OFF duty cycle of intermittent harvesting. Therefore, while the study provides a strong converter-based reference for autonomous MFC energy management, it does not investigate recovery timing as an independent variable for exploiting intermittency-induced improvements in MFC performance.

These observations show a tradeoff between electrochemical relevance and implementation relevance. Capacitor-based and cell-voltage-controlled approaches reveal that intermittent operation can be beneficial, but in both cases the resulting frequency and duty cycle remain strongly influenced by the instantaneous state of the PMFC. As a consequence, the imposed timing and the evolving biological state of the cell become difficult to disentangle, complicating both optimization and scientific interpretation of the observed power variations. Converter-based architectures, in contrast, are the most relevant for embedded applications, yet they do not provide explicit control over the recovery timing needed to systematically study the effect of intermittency in a real battery-free setup.

From the three-electrode perspective discussed in Sec. 2.1.1, independent timing control is essential to separate the imposed electrical perturbation from the subsequent electrode and biofilm response. Although a two-electrode harvester cannot independently control electrode potentials, it can preserve explicit control of polarization time, recovery time and duty cycle. This translates the controlled-perturbation principle of intermittent polarization into a practical harvesting architecture, enabling its effects on harvested power and PMFC electrochemical state to be studied separately.

This motivates the approach pursued in this work: a practical intermittent-harvesting architecture for PMFC-powered embedded systems should combine the implementation relevance of converter-based approaches with an explicit control of frequency and duty cycle, so as to decouple the imposed intermittency pattern from the instantaneous PMFC state and allow the effect of recovery timing to be both exploited and systematically studied.

2.1.3 Toward scalable and observable intermittent PMFC operation

A parallel line of work explores how intermittent operation can be extended beyond a single cell, either by interleaving extraction from multiple cells or by directly aggregating them through parallel/series connection. Quaglio *et al.* [65] proposed alternating between two anodes in a single MFC, thereby preserving the benefits of intermittency while providing a continuous power input. Kim *et al.* [66] and Costilla-Reyes *et al.* [50] investigated multi-cell strategies based on interleaved extraction. These studies highlight both the opportunity and the difficulty of scaling intermittent harvesting. On the one hand, interleaving can smoothen the power supply while granting recovery time to individual PMFCs; on the other hand, it constrains the duty cycle experienced by each cell, potentially preventing operation at the duty cycle that would maximize its intermittent power production. Therefore, once suitable frequency and duty-cycle conditions have been identified at the single-cell level, they can be transferred to multi-cell and multi-anode interleaved architectures, provided that the interleaving scheme preserves the recovery time required by each unit. PMFC aggregation through direct series or parallel connection has also been investigated as a power scale-up strategy [67, 68]. However, aggregation introduces a different set of challenges: cell-to-cell mismatch in voltage and current production due to different state or recovery dynamics lead to both internal PMFC array loading [69], whereby stronger cells partially load weaker ones, and even voltage reversal [44, 45]. For these reasons, this thesis develops in the direction of intermittent interleaving rather than direct PMFC aggregation. In this work, we implement both single-cell and dual-cell interleaved control as representative cases, but the proposed timing element can be extended to broader interleaving strategies involving larger PMFC arrays. Although the focus here is on interleaved operation, the same timing approach is not restricted to individual PMFCs: it can also be applied to series- or parallel-connected groups of PMFCs, provided that each group is treated as an independently timed harvesting unit.

Intermittent operation also intersects with the problem of maximum-power/ maximum-energy point tracking. As previously mentioned, several MFC studies have noted that the optimal working point

drifts over time because of microbial evolution, substrate depletion, and evolving electrode conditions [19, 70, 58]. This observation is consistent with the broader view of the PMFC as a source with long transients and strong loading-history dependence. Continuous-loading MPPT for MFCs has been investigated through several approaches, including hysteresis-controller-based MPPT with digitally controlled operating thresholds [19], MPPT with voltage-overshoot avoidance for safer MFC operation [71], net-power-positive MPPT architectures that account for controller consumption [70], and recent multi-cell converter-based PMSs implementing independent MPPT for each MFC [58, 68]. These works show that operating-point control is increasingly mature for continuous MFC harvesting. However, intermittent harvesting introduces an additional control dimension. The target is no longer only the instantaneous operating voltage or equivalent input impedance, but the combination of harvesting time, recovery time, duty cycle, frequency, and cell operating point. Previous intermittent-harvesting studies have optimized parts of this space: Grondin *et al.* [13] investigated duty cycle and cell-voltage thresholds, while Liang *et al.* [62] jointly studied frequency, duty cycle, and external resistance under transient-state regulation. Nevertheless, in these works the operating point is imposed through external loads or voltage thresholds, rather than through an ultra-low-power input-modulated converter integrated into an autonomous energy-harvesting node. Therefore, although previous studies have addressed continuous-loading MPPT or partial optimization of intermittent operation, the joint optimization of intermittent timing and converter-imposed operating voltage remains largely unexplored, especially in two-electrode PMFC harvesting architectures aimed at embedded-system powering. Even though this thesis does not directly address the identification of the optimal maximum-power or maximum-energy operating point, it lays the groundwork for an ultra-low-power architecture capable of both implementing intermittent operation and controlling the PMFC cell voltage. This is achieved by externally imposing the cell operating point through the input-modulation capability of the energy-harvesting converter; beyond duty-cycle and frequency control, the proposed architecture can be extended to future PMFC-specific MPPT/MEPT strategies and to cell-voltage control methods.

This thesis also focuses on the integration of ultra-low-power

PMFC state observation. The state of a PMFC determines both its instantaneous power output and its long-term energy-delivery capability, but the relevant variables are difficult to access directly. In this thesis, *PMFC state* is used operationally to describe the evolving electrochemical condition of the cell as reflected by externally measurable quantities such as cell voltage, electrode potentials, and impedance spectra; it does not imply identification of a unique internal biological or electrochemical state. Accordingly, *state observation* refers to tracking changes in these diagnostic signatures over time rather than estimating a complete set of internal state variables. Assessing the biofilm, electrode surface, substrate, or plant-root interface typically requires destructive sampling, which alters the PMFC and prevents continuous in-situ monitoring. Indirect electrochemical measurements are therefore attractive because they can probe the cell through its own electrodes. In this framework, electrochemical observation complements two-electrode harvesting control: the harvesting interface imposes cell-level variables such as operating voltage and extraction timing, while electrochemical measurements can reveal how the individual electrode response evolves under the resulting loading history. Although this does not reproduce three-electrode control during harvesting, it provides a means to relate practical energy-extraction conditions to electrode-level electrochemical changes. However, these techniques are usually performed with laboratory instrumentation and require non-negligible energy budgets, making their use in autonomous battery-free nodes challenging. Moreover, since most electrochemical techniques control the cell operating point, the measurement itself perturbs the PMFC loading profile and can influence its subsequent behaviour [72].

Such state observation is also relevant for energy management. Although the high internal impedance of PMFCs is a major limitation for power output, it can also provide information on electrochemical, biological, and transport processes inside the cell. Impedance-based models may therefore support PMFC-specific MPPT and MEPT strategies, lifetime optimization, and biosensing. In this perspective, EIS could eventually provide feedback for adapting the harvesting operating point, period, or duty cycle as the PMFC electrochemical state evolves. Nevertheless, PMFC modeling remains difficult: while Randles-type equivalent circuits are widely used, there is no universally accepted model able to capture the behavior

of these living electrochemical systems. Existing approaches include data-driven models relating temperature, pH, and open-circuit voltage during start-up [73], as well as EIS-based physical models for microbial electrochemical cells [74]. A further motivation is the possible link between electrochemical features and plant state. Carmalin *et al.* [75] highlight the lack of studies on this connection, while suggesting that EIS parameters such as solution resistance may depend on root-exudate type, exudation rate, and the associated bacterial community. Other works apply EIS to PMFCs [37, 76], but often without explicitly considering the plant contribution to the cell's electrobiochemical behavior. If this link is established, PMFCs could act not only as energy sources, but also as inexpensive and portable biosensors for plant-state monitoring.

In this thesis, we move electrochemical state observation from laboratory instrumentation toward a battery-free embedded context. We showed in previous work that measurements such as EIS [77] and CV [25] can be implemented in a self-powered PMFC node, enabling repeated and non-destructive observation of the cell state. This opens the way to PMFC-specific MPPT/MEPT strategies, validation of intermittent-harvesting effects, and biosensing applications in which the PMFC operates simultaneously as an energy source and as a probe of the living system. As a proof of concept, we compare the EIS spectra of PMFCs operated under three conditions: continuous harvesting, intermittent harvesting through the proposed architecture, and unloaded operation. This comparison highlights both the effect of the imposed loading profile on the PMFC impedance response and the fact that the EIS measurement itself acts as an electrical perturbation of the cell. We further fit the EIS spectra of continuously and intermittently loaded electrodes to an equivalent-circuit model, showing how the measurement load influences the extracted impedance parameters.

Two broad conclusions emerge from the literature. First, intermittent power draw is a promising strategy for PMFC-powered embedded systems. By alternating harvesting and recovery phases, intermittency can exploit transient power peaks, accommodate the capacitive and history-dependent behavior of the cell, and potentially influence biofilm activity, mass transport, and long-term stability. Therefore, the objective is not only to maximize instantaneous power, but to improve the energy that can be sustainably extracted over time.

Second, the existing literature remains fragmented. Three-electrode intermittent-polarization studies provide valuable insight into biofilm dynamics, charge storage, and electrode conditioning, but they do not directly implement usable energy harvesting. Conversely, two-electrode intermittent-harvesting studies are relevant to deployable PMFC-powered nodes, but often leave the intermittent timing variables tied to capacitor charging, voltage thresholds, or source recovery dynamics. As a result, frequency, duty cycle, recovery time, and operating voltage are rarely controlled independently, making it difficult to separate the effect of the imposed loading strategy from the evolving state of the PMFC itself.

This motivates the approach pursued in this thesis: the development of a battery-free intermittent-harvesting architecture in which the harvesting profile can be explicitly imposed while the PMFC state can be observed electrochemically. By combining ultra-low-power converter-based energy extraction with controlled timing and self-powered EIS/CV measurements, this work moves toward a co-design framework where harvesting, stabilization, and monitoring are treated as coupled aspects of the same problem. In this perspective, intermittent operation is not merely a power-management technique, but a way to actively shape and observe the PMFC as a living electrochemical energy source.

2.1.4 Battery-free EIS capability

The intermittent-harvesting interface presented above addresses the power-extraction side of PMFC operation. The complementary requirement is the ability to observe the PMFC state without relying on laboratory instrumentation. In previous work [20], we addressed this point by designing a battery-free PMFC-powered EIS node in which the same PMFC electrodes are used both for energy extraction and for impedance measurement. The architecture combines a PMFC-tailored energy-harvesting stage, an ultra-low-power EIS analog front end, energy-aware task triggering, and LoRa communication, thereby demonstrating that PMFC state observation can be integrated into an autonomous sensing node rather than performed only with benchtop instrumentation.

The node uses the MAX30134 analog front end for two-terminal EIS measurements. At system level, the PMFC electrodes are shared between the harvesting and measurement paths: during energy ex-

traction they are connected to the harvesting converter, whereas during an EIS acquisition they are switched to the EIS analog front end. The measurement is triggered only after sufficient energy has been accumulated in the storage capacitor; once the sweep is completed, the EIS circuitry is power-gated and the electrodes are returned to the harvesting path. In this way, impedance measurement is integrated as an energy-admissible task within the battery-free operating cycle.

The MAX30134 response was validated against a reference potentiostat. The measurements showed that the integrated front end can reproduce the PMFC impedance spectrum with sufficient fidelity to track differences across cells and over time. In particular, impedance spectra varied appreciably over a three-day observation window and differed between active and inactive PMFCs, supporting the use of EIS as an indirect indicator of PMFC state. The same work also quantified the energy cost of embedded impedance monitoring: a full 19-point sweep from 21.3 mHz to 21.8 kHz required 3.64 J, while the corresponding LoRa transmission required only 7.4 mJ. To put this energy requirement into perspective, 3.64 J corresponds to the usable energy of approximately 0.81 F when the storage voltage decreases from 3.6 V to 2 V; the original implementation therefore adopted a 1 F supercapacitor. For comparison, the intermittently operated PMFCs in Sec. 3.2 reached average powers of 122.4 µW and 67.6 µW during the final experimental interval. At these power levels, accumulating 3.64 J would require approximately 8.3 h and 15 h, respectively, neglecting conversion and storage losses. The EIS sweep therefore dominates the energy budget, but remains compatible with PMFC-level energy production when its execution is conditioned on sufficient stored energy.

In the original implementation, this energy-aware triggering is performed by a programmable voltage monitor, which launches the EIS task once the storage voltage reaches the required threshold. The same execution principle can be transferred directly to the CapDYN framework developed in Chapter 4: rather than relying on a voltage threshold, the profiled EIS energy requirement can be associated with a CapDYN storage stage and the measurement released through the energy-state interface only when that stage is reached.

The previous implementation therefore provides the system-level feasibility basis for battery-free PMFC impedance monitoring. Chapter 3 addresses the complementary question of whether the im-

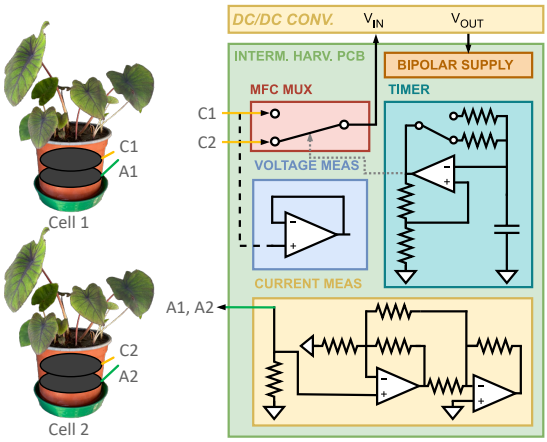


Figure 2.3: Intermittent harvesting architecture, featuring interleaved extraction from two PMFCs, current and voltage measurement front end.

posed loading history produces measurable impedance signatures. For this purpose, potentiostat-based three-electrode EIS is used to resolve the anode and cathode separately at polarization points derived from the electrode potentials reached during two-electrode harvesting. The former study establishes that EIS can be implemented and scheduled within a battery-free PMFC-powered node, while the experimental study in Chapter 3 establishes that loading-history-dependent electrode impedance changes can be detected under harvesting-relevant conditions. Together, these results provide the basis for future embedded EIS systems that autonomously track such changes during operation and potentially use them to adapt the harvesting parameters.

2.2 Intermittent loading design and implementation

The proposed intermittent harvesting architecture (shown in Fig 2.3) functions as an interface between the energy source and the DC–DC converter, thereby integrating into the framework of established energy-harvesting architectures while extending their operational capabilities. The core element of the architecture is the

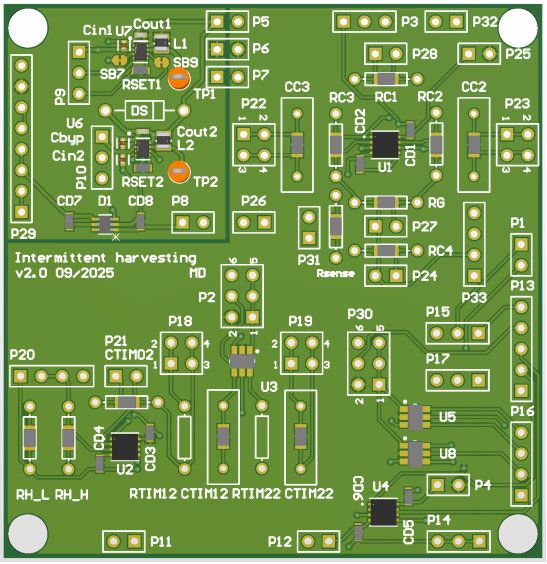


Figure 2.4: PCB implementation of the intermittent-harvesting architecture. The bipolar supply is located in the top-left section of the board. The top-right section implements the current-sensing front end, whereas the bottom-right section implements the voltage-sensing front end. The timer element is located in the bottom-left section.

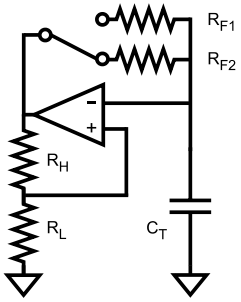


Figure 2.5: Schmitt trigger relaxation oscillator used as timer for the proposed intermittent harvesting interface.

timer, which determines the frequency and duty cycle with which each PMFC actively supplies power. During its active interval, the selected PMFC is connected to the input of the DC–DC converter through the input multiplexer. The interface also integrates an ultra-low power voltage and current front end for monitoring the PMFC power production and is powered through a bipolar supply.

In view of the gaps identified in the related work, we believe that the proposed interface provides a useful experimental and technological platform for advancing intermittent PMFC harvesting. By decoupling the imposed harvesting frequency and duty cycle from the instantaneous recovery dynamics of the cell, while preserving compatibility with converter-based embedded power delivery, the architecture addresses the need for a practical two-electrode implementation in which intermittency can be both exploited and systematically studied. This explicit timing control enables the controlled-study methodology of three-electrode intermittent polarization to be applied to practical two-electrode harvesting. Together with its ultra-low-power implementation tailored to battery-free systems and its integrated voltage/current monitoring capability, this interface can therefore support future research on intermittent harvesting, PMFC-specific energy-management strategies, and the use of PMFCs in autonomous nodes.

2.2.1 Timing and input multiplexer

The decision to rely on an analog timer was motivated by two main considerations. First, the use of an always-on MCU would be in-

consistent with the scope of this thesis, which targets battery-free intermittent-computing systems and therefore requires the quiescent consumption of the control circuitry to be minimized. Second, we envision that, if intermittent harvesting becomes an established operating mode for PMFC-powered systems, the timing function could be integrated directly into energy-harvesting DC–DC converter chips. The external timer adopted in this prototype can therefore be interpreted as a discrete implementation of a control block that could eventually be embedded in PMFC-specific power-management integrated circuits, replacing control strategies designed for faster energy sources and not optimized for the slow, history-dependent dynamics of PMFCs [9].

The timer circuit is implemented as a Schmitt-trigger relaxation oscillator. The operational amplifier is used as a comparator with positive feedback, while the timing capacitor C_T is periodically charged and discharged through the selectable feedback resistors R_{F1} and R_{F2}. Since the operational amplifier is supplied from a bipolar voltage source, its output can be approximated as switching symmetrically between the two saturation levels

$$V_{OH} = +V_{sat}, \qquad V_{OL} = -V_{sat}. \tag{2.1}$$

The non-inverting input is biased by the resistive divider composed of R_H and R_L, which is connected between the output of the comparator and ground. Therefore, the voltage at the non-inverting input is a fixed fraction of the output voltage. Defining

$$\beta = \frac{R_L}{R_H + R_L}, \tag{2.2}$$

the Schmitt-trigger thresholds are

$$V_{TH,+} = +\beta V_{sat}, \qquad V_{TH,-} = -\beta V_{sat}. \tag{2.3}$$

The voltage across the timing capacitor C_T, connected to the inverting input, therefore oscillates between these two thresholds. The resulting capacitor voltage swing is

$$\Delta V_{C_T} = V_{TH,+} - V_{TH,-} = 2\beta V_{sat}. \tag{2.4}$$

When the comparator output is high, the capacitor charges toward $+V_{sat}$ through R_{F1}. The high-output interval corresponds to

the time required for V_{C_T} to increase from $V_{TH,-}$ to $V_{TH,+}$. Therefore,

$$t_{HIGH} = R_{F1} C_T \ln\left(\frac{V_{sat} - V_{TH,-}}{V_{sat} - V_{TH,+}}\right). \tag{2.5}$$

Substituting the threshold expressions gives

$$t_{HIGH} = R_{F1} C_T \ln\left(\frac{1+\beta}{1-\beta}\right) \tag{2.6}$$

When the comparator output is low, the capacitor discharges toward $-V_{sat}$ through R_{F2}. The low-output interval corresponds to the time required for V_{C_T} to decrease from $V_{TH,+}$ to $V_{TH,-}$, yielding

$$t_{LOW} = R_{F2} C_T \ln\left(\frac{V_{TH,+} + V_{sat}}{V_{TH,-} + V_{sat}}\right) \tag{2.7}$$

After substitution,

$$t_{LOW} = R_{F2} C_T \ln\left(\frac{1+\beta}{1-\beta}\right) \tag{2.8}$$

The oscillation period is therefore:

$$T = t_{HIGH} + t_{LOW} = (R_{F1} + R_{F2})\, C_T \ln\left(\frac{1+\beta}{1-\beta}\right) \tag{2.9}$$

and the oscillation frequency is

$$f = \frac{1}{(R_{F1} + R_{F2})\, C_T \ln\left(\frac{1+\beta}{1-\beta}\right)} \tag{2.10}$$

Finally, the duty cycle is defined as $D = \frac{t_{HIGH}}{T}$. Under the bipolar-supply approximation, the logarithmic term is identical during the high and low phases. As a result, the duty cycle depends only on the ratio between the two timing resistors:

$$D = \frac{R_{F1}}{R_{F1} + R_{F2}} \tag{2.11}$$

Eq. 2.11 shows that R_{F1} and R_{F2} independently set the duty cycle of the oscillation, while Eq. 2.9 shows that C_T, $R_{F1} + R_{F2}$ and the Schmitt-trigger factor β scale the overall oscillation period. In particular, increasing C_T or either timing resistance increases

the oscillation period, whereas increasing β increases the capacitor voltage swing and also increases the period through the logarithmic factor.

Single-cell intermittent harvesting necessarily includes intervals during which no input power is delivered to the DC–DC converter. The use of multiple PMFCs can mitigate this limitation, since the input power flow to the converter can be maintained continuously while each individual cell is still operated intermittently. In this way, interleaving enables recovery intervals at the cell level without introducing intervals in which no power is available at the converter input. Because the timer output is binary, the proposed architecture can directly use this signal to interleave the extraction from up to two PMFCs. The corresponding switching stage can be implemented with a 2:1 multiplexer/SPDT analog switch, with the two cells operated at complementary duty cycles.

Extending the architecture to more than two cells requires an additional control stage, since the binary timer signal can no longer directly address all the input-multiplexer states. In this case, the timer signal can be used as a clock for a counter, which updates the selected PMFC at each active timer edge. The timer duty cycle then no longer determines the duty cycle experienced by each PMFC; instead, assuming uniform interleaving and one active cell at a time, the duty cycle of each cell is set by the total number of cells:

$$D_i = \frac{1}{N_{\mathrm{CELLS}}}. \tag{2.12}$$

Thus, increasing the number of interleaved PMFCs improves the continuity of the power flow at the converter input, but also constrains the harvesting duty cycle assigned to each individual cell. In addition, scaling the input multiplexer to a larger number of channels may increase the series resistance, leakage paths, and parasitic contributions introduced between the PMFCs and the DC–DC converter input, depending on the selected switching technology.

The current implementation of the intermittent-harvesting PCB is shown in Fig. 2.4. The board features two PMFC inputs, which were selected to enable both continuous power delivery to the converter input and flexible exploration of the harvesting duty cycle. Scaling the architecture to a larger number of cells, as well as identifying the optimum number of interleaved PMFCs, represents a natural continuation of this work once suitable duty-cycle conditions

have been experimentally consolidated using the proposed platform.

The voltage- and current-sensing front ends, as well as the timer's Schmitt trigger, were implemented using Maxim Integrated/Analog Devices MAX40018 operational amplifiers, consuming only 400 nA per channel. The switching stage was implemented using an Analog Devices ADG819 SPDT analog switch, directly controlled from the timer signal. If expansion to a larger number of PMFCs is required, the timer output can be used as the clock input of a low-power CMOS binary counter, such as the CD4520B. For four PMFCs, the two least-significant counter bits can drive a 4:1 multiplexer such as the ADG704. For eight PMFCs, the three least-significant bits can drive an ADG708, while for sixteen PMFCs all four counter bits can drive an ADG706. Finally, up to thirty-two PMFCs can be interleaved by cascading the two 4-bit counters available within a dual CD4520B package and using the five least-significant bits of the resulting counter to address a 32:1 multiplexer such as the ADG732. The bipolar supply was implemented using two Texas Instruments 60-nA quiescent current TPS62840 converters, one operated in the standard buck configuration and the other in an inverting buck configuration, following the topology described in the corresponding Texas Instruments application report.[1] The intermittent-harvesting interface was finally connected to a Texas Instruments BQ25505 energy-harvesting DC–DC converter, which provides an established reference platform for ultra-low-power extraction from high-impedance energy sources.

According to the literature, the relevant timing-period range for intermittent PMFC operation extends from milliseconds [63] seconds to several hours [46]. While the lower end of this range was readily achievable with the proposed timer, reaching the upper end required a more careful design. Since the maximum practical value of low-leakage capacitors is limited, longer periods were obtained by increasing the feedback resistance. This reduces the charge and discharge currents of the timing capacitor and, consequently, makes the timer more sensitive to parasitic leakage currents. At long timing periods, the timer accuracy is therefore limited by the leakage paths associated with the timing network. These include the leakage of the capacitor itself, the leakage of the analog switch used for duty-cycle control, and the input bias current of the timer op-amp. The MAX40018 was selected to minimize the latter contribution,

[1] `https://www.ti.com/lit/an/slvafu9/slvafu9.pdf`

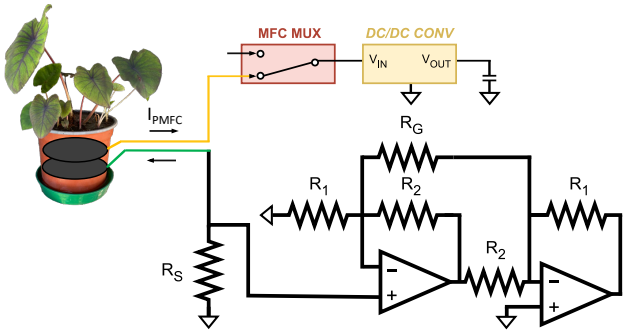


Figure 2.6: Current measurement front end schematic.

as it features an input bias current of 0.1 pA, while ceramic capacitors were used to reduce dielectric leakage. Under these conditions, the dominant parasitic contribution at very low timing currents was found to be the leakage of the selected analog switch, the ADG819, whose source-off and channel-on leakage currents are specified in the order of 10 pA. The effect of the timer feedback current nearing the ADG819 leakage is particularly evident in the duty-cycle accuracy. Since the two states of the SPDT switch introduce different leakage paths, the effective charge and discharge currents of the timing capacitor become asymmetric. As a result, the measured duty cycle deviates from its expected value, even when the nominal timing network is configured for the desired duty cycle. In conclusion, the timer feedback current should remain sufficiently larger than the dominant leakage currents in the timing network to achieve accurate timing and duty-cycle control.

2.2.2 Current and voltage measurement front end

Voltage sensing is performed by unity-gain buffering of the PMFC cell voltages. The cell voltages are monitored during both active and inactive phases, in order to verify the imposed cell polarization during harvesting and the voltage recovery during open-circuit intervals, respectively.

Current sensing is performed only on the active PMFC, since it is the only cell delivering current to the DC–DC converter. The selected current-sensing topology, shown in Fig. 2.6, is based on a low-side shunt resistor whose voltage drop is amplified by a two-op-

amp instrumentation-amplifier stage. This configuration provides high input impedance while limiting the number of active components, thereby reducing the power consumption of the sensing front end. Since the shunt resistor is connected to ground, the resulting shunt voltage is single-ended:

$$V_{\mathrm{SH}} = I_{\mathrm{PMFC}} R_S. \tag{2.13}$$

To derive the gain, the resistors are first considered non-identical. Let R_{1a} and R_{2a} be the input and feedback resistors of the first stage, and let R_{2b} and R_{1b} be the input and feedback resistors of the inverting second stage. Assuming ideal op-amps, the inverting input of the first op-amp is at V_{SH}, while the inverting input of the second op-amp is at virtual ground. Applying KCL at the first inverting node gives

$$V_1 = V_{\mathrm{SH}} \left(1 + \frac{R_{2a}}{R_{1a}} + \frac{R_{2a}}{R_G}\right), \tag{2.14}$$

where V_1 is the output of the first op-amp. Applying KCL at the virtual-ground node of the second op-amp gives

$$V_{\mathrm{OUT}} = -R_{1b} \left(\frac{V_1}{R_{2b}} + \frac{V_{\mathrm{SH}}}{R_G}\right). \tag{2.15}$$

Therefore,

$$V_{\mathrm{OUT}} = -V_{\mathrm{SH}} \left[\frac{R_{1b}}{R_{2b}} \left(1 + \frac{R_{2a}}{R_{1a}} + \frac{R_{2a}}{R_G}\right) + \frac{R_{1b}}{R_G}\right]. \tag{2.16}$$

If the two resistor ratios are matched, namely

$$\frac{R_{1a}}{R_{2a}} = \frac{R_{1b}}{R_{2b}}, \tag{2.17}$$

the expression simplifies to

$$V_{\mathrm{OUT}} = -V_{\mathrm{SH}} \left(1 + \frac{R_{1b}}{R_{2b}} + \frac{R_{1a} + R_{1b}}{R_G}\right). \tag{2.18}$$

If, in addition, the corresponding resistors are chosen equal,

$$R_{1a} = R_{1b} = R_1, \qquad R_{2a} = R_{2b} = R_2, \tag{2.19}$$

the output voltage becomes

$$V_{\mathrm{OUT}} = -V_{\mathrm{SH}} \left(1 + \frac{R_1}{R_2} + 2\frac{R_1}{R_G}\right). \tag{2.20}$$

Finally, substituting $V_{\mathrm{SH}} = I_{\mathrm{PMFC}} R_S$,

$$V_{\mathrm{OUT}} = -I_{\mathrm{PMFC}} R_S G_I = -I_{\mathrm{PMFC}} R_S \left(1 + \frac{R_1}{R_2} + 2\frac{R_1}{R_G}\right) \tag{2.21}$$

Noise analysis of the current-sensing front end

A preliminary noise analysis was carried out to identify the dominant noise contributions of the current-sensing front end and to guide the choice of the shunt resistance and amplifier gain. The analysis combined a first-order output-noise model with empirical measurements in which one design parameter was varied at a time. The tested parameters were the absolute value of the feedback resistors, the shunt resistance, the capacitance placed in parallel with the shunt, and the current-sense gain. The capacitor placed across the shunt was used to assess whether the noise associated with the measured current was dominated by high-frequency components.

Since the measured quantity is the output voltage spectral density, the different noise contributions are combined in terms of power spectral density. The output-noise power spectral density can be written as

$$\begin{aligned} S_{v,o}(f) = &\sum_{j=1}^{N_{\mathrm{op}}} |H_{e_j}(f)|^2 e_{n,j}^2(f) + \sum_k |H_{R_k}(f)|^2 4kTR_k \\ &+ \sum_{j=1}^{N_{\mathrm{op}}} \left(|H_{i_{j,+}}(f)|^2 i_{n,j,+}^2(f) + |H_{i_{j,-}}(f)|^2 i_{n,j,-}^2(f)\right) \\ &+ |H_I(f)|^2 i_{n,\mathrm{meas}}^2(f). \end{aligned} \tag{2.22}$$

Here, $e_{n,j}$ denotes the input voltage-noise density of the j-th operational amplifier, $4kTR_k$ is the thermal-noise power spectral density of the k-th resistor, and $i_{n,j,+}$ and $i_{n,j,-}$ are the input current-noise densities of the non-inverting and inverting input terminals of the j-th operational amplifier. The transfer functions $H_{e_j}(f)$, $H_{R_k}(f)$, $H_{i_{j,+}}(f)$, and $H_{i_{j,-}}(f)$ refer each contribution to the output. The term $i_{n,\mathrm{meas}}$ represents fluctuations of the measured PMFC current,

converted into voltage by the shunt resistor and then amplified by the current-sense stage. The measurement-chain floor was not explicitly subtracted, but was treated as a constant contribution across configurations because all acquisitions used the same setup and processing.

For the matched-resistor implementation, the low-frequency signal gain of the current-sensing stage is specified by G_I in Eq. 2.21. For the baseline configuration, $R_1 = 400$ kΩ, $R_2 = 45$ kΩ, and $R_G = 2.8$ kΩ result in $G_I \simeq 295.6$.

The two voltage-noise transfer functions have the same magnitude as the signal gain, G_I. Therefore,

$$e_{o,\mathrm{op}}(f) \simeq \sqrt{2}\, G_I e_n(f), \tag{2.23}$$

where the factor $\sqrt{2}$ accounts for the fact that the two uncorrelated voltage-noise sources add in power.

An order-of-magnitude noise budget was then evaluated over the 0.1–5 Hz band used for the empirical comparison, corresponding to $B = 4.9$ Hz. The MAX40018 datasheet specifies a low-frequency input voltage noise of 7 $\mu\mathrm{V_{RMS}}$ from 0.1 to 10 Hz. Using this value as a conservative estimate gives a combined input-referred contribution of approximately

$$e_{\mathrm{op,in}} \simeq \sqrt{2} \cdot 7 \simeq 9.9\ \mu\mathrm{V_{RMS}}, \tag{2.24}$$

or, referred to the output,

$$e_{\mathrm{op,out}} \simeq G_I e_{\mathrm{op,in}} \simeq 295.6 \cdot 9.9\ \mu\mathrm{V} \simeq 2.9\ \mathrm{mV_{RMS}}. \tag{2.25}$$

This value should be interpreted as a conservative upper-bound estimate, since the datasheet value is specified over 0.1–10 Hz rather than exactly over the 0.1–5 Hz band used here. Nevertheless, it shows that the op-amp input voltage noise is expected to be the largest contribution in the first-order budget.

The thermal noise of the resistive network is much smaller. For the baseline configuration, propagating the resistor thermal-noise sources through the two-op-amp network gives an output-referred contribution of approximately

$$e_{o,R,\mathrm{rms}} \simeq 4.4\ \mu\mathrm{V_{RMS}} \tag{2.26}$$

over the 0.1–5 Hz band. If all resistor values are increased by a factor of ten while preserving the same gain, this contribution increases by

approximately $\sqrt{10}$, reaching about

$$e_{o,R,\mathrm{rms}} \simeq 14\ \mu\mathrm{V_{RMS}}. \tag{2.27}$$

Both values are far below the expected op-amp voltage-noise contribution, which explains why the 400 kΩ and 4 MΩ implementations exhibit very similar measured spectra.

The intrinsic thermal noise of the shunt resistor is also negligible. Its voltage-noise density is

$$e_{n,R_S} = \sqrt{4kTR_S}. \tag{2.28}$$

At room temperature, this gives approximately 0.13 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ for $R_S = 1\ \Omega$, 0.41 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ for $R_S = 10\ \Omega$, and 1.3 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ for $R_S = 100\ \Omega$. Even after amplification by $G_I \simeq 295.6$, the corresponding RMS output contributions over 0.1–5 Hz remain below 1 $\mu\mathrm{V_{RMS}}$. Therefore, the shunt thermal noise cannot explain the measured output-noise levels or the weak dependence on shunt resistance.

Since the MAX40018 datasheet does not specify the input current-noise density, this term was estimated only for sensitivity analysis. A conservative value was chosen from the shot-noise estimate associated with the input bias current,

$$i_n \simeq \sqrt{2qI_B}. \tag{2.29}$$

This relation is commonly used to estimate the shot-noise component of op-amp input current noise when the current noise is dominated by the input bias current. Using the typical MAX40018 input bias current (0.1 pA) gives a sub-$\mathrm{fA}/\sqrt{\mathrm{Hz}}$ value, while using the worst-case bias current over temperature (200 pA) gives a value of a few $\mathrm{fA}/\sqrt{\mathrm{Hz}}$. Therefore, 10 $\mathrm{fA}/\sqrt{\mathrm{Hz}}$ was adopted as a conservative round-number estimate, not as a datasheet specification. If this current noise flows through an impedance on the order of 400 kΩ, the corresponding input-referred RMS voltage noise over 0.1–5 Hz is

$$e_{i,\mathrm{in}} \simeq i_{n,\mathrm{op}} R_1 \sqrt{B} = 10\ \mathrm{fA}/\sqrt{\mathrm{Hz}} \cdot 400\ \mathrm{k\Omega} \cdot \sqrt{4.9} \simeq 9\ \mathrm{nV_{RMS}}. \tag{2.30}$$

Referred to the output, this gives approximately

$$e_{i,\mathrm{out}} \simeq G_I e_{i,\mathrm{in}} \simeq 2.6\ \mu\mathrm{V_{RMS}}. \tag{2.31}$$

Even for the 4 MΩ implementation, the same estimate gives an output-referred contribution of only about 26 μV_{RMS}. The op-amp input current-noise contribution is therefore well below the voltage-noise contribution and was considered negligible in the first-order noise budget.

The empirical comparisons were quantified by integrating the measured output voltage spectral density over the frequency band of interest. The spectra were acquired with PicoScope in spectrum mode and exported in dBu. The exported values were converted to voltage spectral density as

$$e_x(f_i) = 0.775 \cdot 10^{L_x(f_i)/20}, \tag{2.32}$$

where $L_x(f_i)$ is the exported level in dBu for channel x, and $e_x(f_i)$ is expressed in $V/\sqrt{Hz}$. The equivalent RMS noise voltage over the selected interval $[f_1, f_2]$ was computed as

$$N_x = \sqrt{\sum_{f_i \in [f_1, f_2]} e_x^2(f_i) \Delta f_i}. \tag{2.33}$$

The integration interval was set to 0.1–5 Hz, in order to focus on the low-frequency region relevant to PMFC current monitoring while remaining within the closed-loop bandwidth of all tested gain configurations. For each configuration, five independent acquisitions were processed, and the reported value corresponds to the mean and standard deviation of the integrated noise.

The empirical tests were all performed starting from a baseline configuration defined by $R_S = 1\ \Omega$, no capacitor in parallel with the shunt, $R_1 = 400$ kΩ, $R_2 = 45$ kΩ, and $R_G = 2.8$ kΩ, corresponding to a nominal gain $G_I \simeq 295.6$. The feedback-resistor scaling test preserved the same resistor ratios and therefore the same nominal gain, but increased all feedback-network resistors by a factor of ten, using $R_1 = 4$ MΩ, $R_2 = 450$ kΩ, and $R_G = 28$ kΩ, while keeping $R_S = 1\ \Omega$ and no shunt capacitor. The shunt-resistance sweep was then performed with the baseline amplifier configuration and no shunt capacitor, varying R_S between 1 Ω, 10 Ω, and 100 Ω. The capacitor sweep used the baseline amplifier configuration and $R_S = 1\ \Omega$, while adding 100 μF, 2.2 mF, and 50 mF in parallel with the shunt. Finally, the gain sweep was performed with $R_S = 1\ \Omega$, no shunt capacitor, $R_1 = 400$ kΩ, and $R_G = 2.8$ kΩ, while reducing

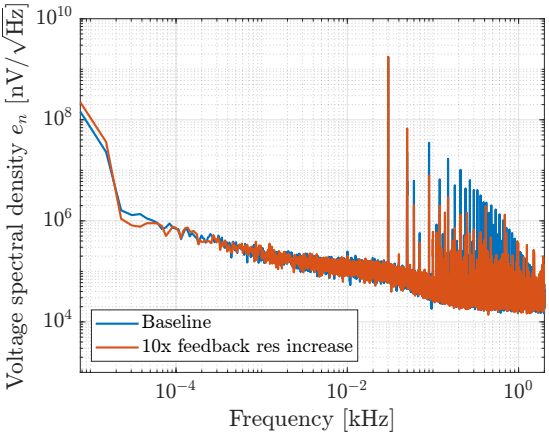


Figure 2.7: Feedback resistor scaling effect on output voltage spectral density of the current sense amplifier.

R_2 from 46 kΩ to 4.6 kΩ and 460 Ω. This increased the gain from approximately 295 to 375 and 1175, respectively.

Feedback resistor scaling: the first empirical test evaluated the effect of increasing the absolute value of the resistors in the feedback network while preserving the same nominal gain. If the feedback-network thermal noise were dominant, increasing all resistor values by a factor of ten would be expected to increase the corresponding voltage-noise contribution by approximately $\sqrt{10} \simeq 3.16$. Instead, the measured output voltage integrated noise changed only weakly, as shown in Fig. 2.7. Under power-supply operation, the noise changed from 529.9 ± 13.4 μV_{RMS} for the baseline implementation to 501.8 ± 16.5 μV_{RMS} for the 10x feedback resistance increase implementation, corresponding to a ratio of 0.95. Under battery operation, the noise changed from 525.8 ± 17.6 μV_{RMS} for the baseline implementation to 526.1 ± 17.1 μV_{RMS} for the 10x implementation, corresponding to a ratio of 1.00. These ratios are far below the $\sqrt{10}$ increase expected from resistor-thermal-noise dominance, indicating that the final output noise is largely unaffected by the thermal noise of the feedback network. This result has secondary implications for ultra-low-power operation: increasing all feedback resistors by a factor of ten preserves the same gain, but reduces the current flowing through the feedback network by the same factor. Thus, the feedback-network power consumption is reduced by

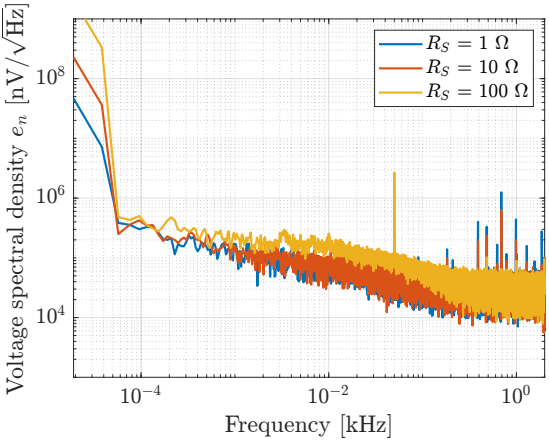


Figure 2.8: Shunt resistance effect on output voltage spectral density of the current sense amplifier.

10×, corresponding to a 90% reduction. This reduction applies to the passive feedback network only; the total front-end consumption also includes the quiescent current of the operational amplifiers.

Shunt resistance sweep: the shunt-resistance sweep showed only a weak dependence of the integrated output noise on R_S (Fig. 2.8). For $R_S = 1\ \Omega$, $10\ \Omega$, and $100\ \Omega$, the measured output integrated noises were $284.3 \pm 9.2\ \mu\text{V}_{\text{RMS}}$, $301.9 \pm 18.6\ \mu\text{V}_{\text{RMS}}$, and $489.1 \pm 271.7\ \mu\text{V}_{\text{RMS}}$, respectively. The large standard deviation of the $100\ \Omega$ case is due to two outlying acquisitions; excluding them gives an average value of approximately $294.7\ \mu\text{V}_{\text{RMS}}$. To quantify the dependence of the integrated noise from the various noise sources, the measured noise ratios were compared with the ratios expected from the candidate dominant-noise mechanisms. Using the $1\ \Omega$ case as reference, the measured ratio for $10\ \Omega$ is

$$\frac{N_o(10\ \Omega)}{N_o(1\ \Omega)} = \frac{301.9}{284.3} \simeq 1.06. \tag{2.34}$$

Using the raw $100\ \Omega$ mean gives

$$\frac{N_o(100\ \Omega)}{N_o(1\ \Omega)} = \frac{489.1}{284.3} \simeq 1.72, \tag{2.35}$$

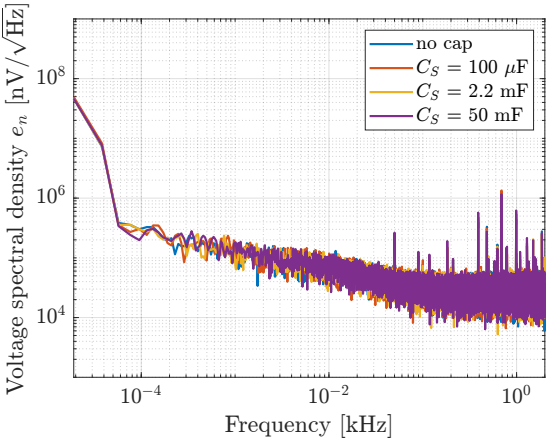


Figure 2.9: Shunt capacitance effect on output voltage spectral density of the current sense amplifier.

whereas excluding the two outlying acquisitions gives

$$\frac{N_o(100\ \Omega)}{N_o(1\ \Omega)} = \frac{294.7}{284.3} \simeq 1.04. \tag{2.36}$$

If the dominant contribution were current fluctuations converted into voltage by the shunt, the expected ratios would be 10 and 100 for the 10 Ω and 100 Ω cases, respectively. If the dominant contribution were the intrinsic thermal noise of the shunt resistor, the expected ratios would instead be $\sqrt{10} \simeq 3.16$ and 10. The measured ratios are much smaller than both predictions. Therefore, the integrated output noise is largely independent of the shunt resistance, indicating that shunt-dependent noise contributions are not dominant.

Shunt capacitance sweep: Capacitors of 100 μF, 2.2 mF, and 50 mF were placed in parallel with the 1 Ω shunt resistor. The parallel R_S-C_S network acts as a first-order low-pass filter for current fluctuations converted into voltage across the shunt, with cutoff frequency

$$f_c = \frac{1}{2\pi R_S C_S}. \tag{2.37}$$

For $R_S = 1\ \Omega$, the corresponding cutoff frequencies are

$$\begin{aligned} f_c(100\ \mu\text{F}) &\simeq 1.59\ \text{kHz}, \\ f_c(2.2\ \text{mF}) &\simeq 72.3\ \text{Hz}, \\ f_c(50\ \text{mF}) &\simeq 3.18\ \text{Hz}. \end{aligned} \tag{2.38}$$

Therefore, the largest capacitor should significantly attenuate current-noise components above a few hertz, while the smaller capacitors mainly affect higher-frequency components. Test results are shown in Fig. 2.9. The corresponding output voltage integrated noise values were $297.7 \pm 22.8\ \mu\text{V}_{\text{RMS}}$, $278.6 \pm 20.3\ \mu\text{V}_{\text{RMS}}$, and $301.1 \pm 16.1\ \mu\text{V}_{\text{RMS}}$. Relative to the no-capacitor case, these correspond to ratios of

$$\frac{297.7}{284.3} \simeq 1.05, \qquad \frac{278.6}{284.3} \simeq 0.98, \qquad \frac{301.1}{284.3} \simeq 1.06. \tag{2.39}$$

Thus, increasing the capacitance does not produce a monotonic reduction in output noise. This behavior is expected for the two smaller capacitor values, whose cutoff frequencies lie above the integration bandwidth and therefore only affect noise components outside the analyzed range. Conversely, if input current noise provided a significant contribution within the measured band, the largest capacitor would be expected to reduce the integrated noise because its cutoff frequency falls within the integration bandwidth. The absence of this reduction indicates that the dominant contribution to output noise is not measured current noise, which would otherwise be filtered by the input shunt capacitance.

Gain sweep: the gain sweep provides the clearest trend, as shown in Fig. 2.10. Reducing R_2 from 46 kΩ to 4.6 kΩ and 460 Ω increased the gain from approximately 295 to 375 and 1175, respectively. The corresponding output integrated noise values were $284.3 \pm 9.2\ \mu\text{V}_{\text{RMS}}$, $325.0 \pm 11.6\ \mu\text{V}_{\text{RMS}}$, and $988.0 \pm 66.5\ \mu\text{V}_{\text{RMS}}$. Using the lowest-gain case as reference, the gain ratios are

$$\frac{375}{295} \simeq 1.27, \qquad \frac{1175}{295} \simeq 3.98, \tag{2.40}$$

while the corresponding measured noise ratios are

$$\frac{325.0}{284.3} \simeq 1.14, \qquad \frac{988.0}{284.3} \simeq 3.48. \tag{2.41}$$

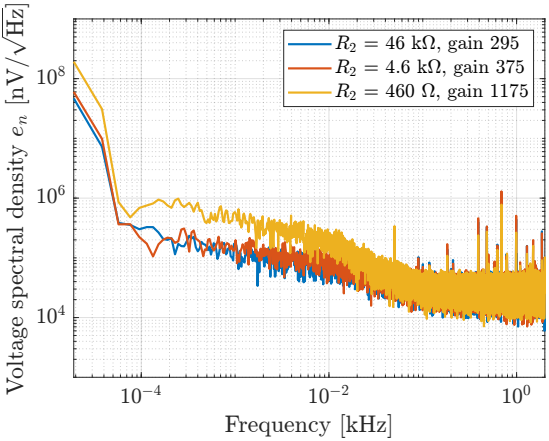


Figure 2.10: Gain effect on output voltage spectral density of the current sense amplifier.

The measured output-noise ratios therefore follow the imposed gain ratios much more closely than they follow the shunt-resistance ratios. This behavior is consistent with a dominant input-referred voltage-noise contribution that is amplified by the current-sense stage.

Table 2.1: Summary of the current-sensing front-end noise analysis in the 0.1–5 Hz band. All measured ratios are normalized to the nominal front-end configuration: original feedback network, $R_S = 1\ \Omega$, no added shunt capacitance, and $G_I = 295$, under the same operating condition.

Test	**Hypothesis tested**	**Expected trend if dominant**	**Measured ratio**	**Conclusion**
Feedback-resistor scaling	Feedback-network thermal noise.	Increasing all feedback resistors by 10× should increase the voltage-noise contribution by $\sqrt{10} \simeq 3.16$.	Ratios are 0.95 under power-supply operation and 1.00 under battery operation.	Feedback network not dominant.
Shunt-resistance sweep	Shunt thermal noise or current fluctuations converted into voltage across R_S.	If current fluctuations dominated, the expected ratios would be 10 and 100. If shunt thermal noise dominated, they would be $\sqrt{10}$ and 10.	Ratios are 1.06 for $R_S = 10\ \Omega$ and 1.72 for $R_S = 100\ \Omega$. Excluding two 100 Ω outliers, the 100 Ω ratio becomes approximately 1.04.	Shunt-dependent noise not dominant.
Shunt-capacitance sweep	Current-noise components filtered by the shunt capacitor.	A sufficiently large shunt capacitor should reduce current-noise components in the integration band.	Ratios are 1.05, 0.98, and 1.06 for 100 μF, 2.2 mF, and 50 mF, respectively.	Current fluctuations not dominant.
Gain sweep	Input-referred amplifier voltage noise.	If input-referred voltage noise dominates, output noise should scale approximately with current-sense gain.	Ratios are 1.14 for $G_I = 375$ and 3.48 for $G_I = 1175$, compared with gain ratios of 1.27 and 3.98.	Dominant noise source.

Overall, the empirical analysis, summarized in Table 2.1, identifies the input-referred voltage noise of the MAX40018 amplifiers as the dominant contribution in the 0.1–5 Hz band. The gain sweep provides the clearest evidence: increasing the gain from 295 to 1175 increases the integrated output noise by a factor of 3.48, close to the imposed gain increase of 3.98. In contrast, increasing the feedback resistors by 10×, increasing the shunt resistance, or adding large capacitors across the shunt does not produce the scaling expected from feedback-resistor thermal noise, shunt thermal noise, or high-frequency current noise.

The feedback-resistor scaling test shows that the passive feedback network is not the noise-limiting element. This is important for ultra-low-power design, because increasing all feedback resistors by 10× preserves the same gain while reducing the current drawn by the feedback network by the same factor. This corresponds to a 10×, or 90%, reduction in passive feedback-network power consumption. The practical upper limit on resistor value is therefore not set by thermal noise in this implementation, but most likely by bias-current-induced errors.

The shunt and capacitor sweeps show that shunt-dependent noise is not dominant. Increasing R_S from 1 Ω to 10 Ω increases the integrated output noise by only 6%, and the 100 Ω case becomes comparable to the baseline once the outlying acquisitions are excluded. Similarly, adding 100 μF, 2.2 mF, and 50 mF across the shunt does not produce a monotonic noise reduction. This indicates that the measured output noise is not dominated by current fluctuations converted into voltage by the shunt.

These results lead to three design guidelines. First, the feedback-resistor values can be increased to reduce passive power consumption, provided that current bias and leakage effects remain acceptable. Second, the current-sense gain should be kept as low as allowed by the acquisition resolution, because increasing gain directly amplifies the dominant input-referred voltage noise and reduces the closed-loop bandwidth according to $f_{\mathrm{BW}} \simeq \mathrm{GBP}/G_I$. Third, the shunt resistor should be made as large as possible to improve signal amplitude and signal-to-noise ratio, but only up to the point where its burden voltage does not significantly perturb the PMFC operating point imposed by the DC–DC converter.

Based on these constraints, the final design uses $R_S = 6.7\ \Omega$, $R_1 = 10\ \mathrm{M}\Omega$, $R_2 = 100\ \mathrm{k}\Omega$, and $R_G = 500\ \mathrm{k}\Omega$, corresponding

to a nominal gain $G_I = 141$. The nominal DC transimpedance is therefore

$$\left.\frac{V_{\text{out}}}{I_{\text{PMFC}}}\right|_{\text{nom}} = G_I R_S = 141 \cdot 6.7\ \Omega \simeq 945\ \Omega, \tag{2.42}$$

corresponding to a nominal output sensitivity of 0.945 mV/μA.

The implemented current-sense conversion was then calibrated experimentally by comparing the current-sense output with the validation current measured by Keithley's DMM7510 digital multimeter. After compensating the time offset and skew between the two acquisitions, the DMM current was interpolated at the DAQ timestamps and a linear fit was performed:

$$I_{\text{meas}} = a_{\text{cal}} V_{\text{out}} + b_{\text{cal}} \tag{2.43}$$

For the 6.7 Ω configuration, the calibration gives

$$I_{\text{meas}} = 1.329 \cdot 10^{-3} V_{\text{out}} + 75.8\ \mu\text{A}, \tag{2.44}$$

with I_{meas} expressed in amperes and V_{out} in volts. The corresponding calibrated transimpedance is

$$R_{T,\text{cal}} = \frac{1}{a_{\text{cal}}} = 752.4\ \Omega, \tag{2.45}$$

corresponding to a calibrated output sensitivity of

$$0.752\ \text{mV}/\mu\text{A} \tag{2.46}$$

The fitted offset corresponds to a zero-current output voltage of

$$V_{\text{out},0} = -\frac{b_{\text{cal}}}{a_{\text{cal}}} \simeq -57\ \text{mV}. \tag{2.47}$$

Assuming a usable output swing of approximately ± 2 V, the calibrated current range is obtained from the fitted conversion:

$$I_{\text{meas,min}} = a_{\text{cal}}(-2\ \text{V}) + b_{\text{cal}} \simeq -2.58\ \text{mA} \tag{2.48}$$

$$I_{\text{meas,max}} = a_{\text{cal}}(2\ \text{V}) + b_{\text{cal}} \simeq 2.73\ \text{mA} \tag{2.49}$$

Since the PMFC harvesting current is expected to be positive during normal operation, the practically relevant positive measurement

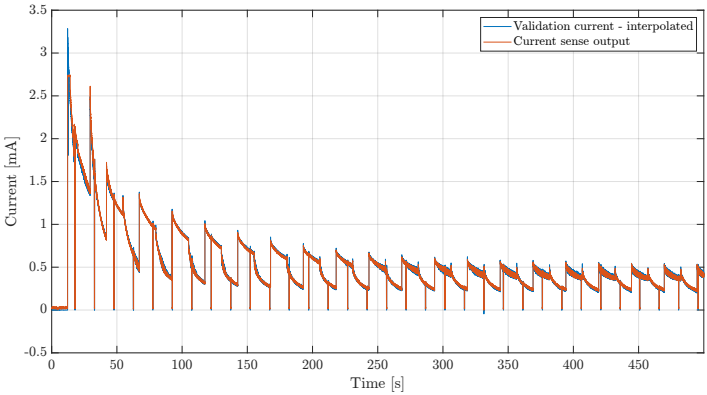


Figure 2.11: Validation of the developed current-sensing front end during intermittent harvesting from two PMFCs. The current estimated from the front end is compared with a reference measurement acquired using a Keithley DMM7510.

range extends up to approximately $I_{\text{meas,max}} \simeq 2.7$ mA. The difference between the nominal and calibrated sensitivities is attributed to component tolerances, offset terms, and non-idealities of the implemented front end.

Figure 2.11 reports the experimental validation of the implemented current-sensing front end during intermittent harvesting from two PMFCs through the proposed interface. The current estimated from the developed sensing circuit is compared with a reference measurement acquired using a Keithley DMM7510.

The comparison shows good agreement between the developed analog front end and the reference instrument. Large current transients are observed immediately after connection, and the first peak may exceed the available measurement range, leading to clipping. However, the regime current values, which are the relevant quantities for estimating the long-term average harvested power, remain within the measurement range of the developed current-sensing infrastructure.

Shunt-burden correction

The calibrated relation in Eq. 2.43 gives the current flowing in the instrumented measurement circuit. This quantity is denoted as I_{meas},

because the current-sensing front end includes the low-side shunt resistor R_S in the PMFC return path. The shunt is therefore both a sensing element and a small series resistance inserted in the PMFC/load loop, so I_{meas} is not the ideal uninstrumented current. In the following, I_0 denotes the estimated current that would flow through the same equivalent external load after removing the shunt, while I_{PMFC} is retained only for generic schematic-level PMFC current definitions.

In the final configuration, $R_S = 6.7\ \Omega$, and the corresponding burden voltage is

$$V_S = I_{\mathrm{meas}} R_S. \tag{2.50}$$

At $I_{\mathrm{meas}} = 1$ mA, the shunt drop is $V_S = 6.7$ mV. Let V_L denote the voltage across the external harvesting interface/load, measured between the PMFC positive terminal and the circuit ground at the top of the low-side shunt. The instantaneous electrode-to-electrode voltage of the PMFC in the instrumented configuration is then

$$V_{\mathrm{PMFC,inst}} = V_L + I_{\mathrm{meas}} R_S. \tag{2.51}$$

The corresponding PMFC-side electrical power in the instrumented configuration is

$$P_{\mathrm{PMFC,inst}} = I_{\mathrm{meas}} V_{\mathrm{PMFC,inst}} = I_{\mathrm{meas}} V_L + I_{\mathrm{meas}}^2 R_S. \tag{2.52}$$

The first term in Eq. 2.52 is the power delivered to the external harvesting interface/load, whereas the second term is the power dissipated in the measurement shunt. Their sum is the PMFC-side electrical power in the instrumented configuration. This accounting includes the shunt voltage drop, but it does not remove the operating-point perturbation caused by the additional series resistance.

The shunt-burden correction was therefore applied in post-processing to the intermittently loaded cells, for which the low-side shunt was part of the current-measurement path. The correction locally linearizes the PMFC response around the measured operating point. For this small perturbation, the cell is represented by a Thevenin-equivalent source with local resistance $R_{\mathrm{int,loc}}$, estimated from the EIS-based analysis described below. The removal of the shunt is then modeled as a small perturbation of the total series resistance, while the equivalent load imposed by the harvesting interface on

the PMFC is assumed unchanged. The resulting current I_0 is therefore a first-order estimate of the no-shunt current, rather than a full model of the nonlinear PMFC behavior.

Under this approximation, the current measured in the instrumented configuration is

$$I_{\text{meas}} = \frac{V_{\text{OC}}}{R_{\text{int}} + R_L + R_S}, \tag{2.53}$$

whereas the current that would flow through the same equivalent external load without the shunt is

$$I_0 = \frac{V_{\text{OC}}}{R_{\text{int}} + R_L}. \tag{2.54}$$

The relative current perturbation introduced by the shunt is therefore

$$\epsilon_I = \frac{I_0 - I_{\text{meas}}}{I_0} = \frac{R_S}{R_{\text{int}} + R_L + R_S}. \tag{2.55}$$

If the load contribution is neglected, setting $R_L = 0$ gives a conservative upper bound. For $R_S = 6.7\ \Omega$ and $R_{\text{int}} = 500$–$100\ \Omega$, this gives

$$\epsilon_I = \frac{6.7}{R_{\text{int}} + 6.7} = 1.32\%\text{–}6.28\%. \tag{2.56}$$

For the corrected power estimates reported in the long-term analysis, R_{int} was not taken from the conservative range 500–100 Ω, but estimated locally from the fitted EIS model of the corresponding cell and measurement period. For each segment, the local PMFC source resistance was computed as

$$R_{\text{int,loc}} = \Re\{Z_A(2\pi f_{\text{EIS}}) + Z_C(2\pi f_{\text{EIS}})\}, \qquad f_{\text{EIS}} = 50\ \text{mHz}, \tag{2.57}$$

where Z_A and Z_C are the fitted anode and cathode impedance functions of the corresponding cell and acquisition segment, obtained from the equivalent-circuit parameters reported in Tabs. 3.3 and 3.4, respectively. The equivalent external load at each active sample was estimated from the measured load voltage and current as

$$R_{L,\text{eq}} = \frac{V_L}{I_{\text{meas}}}. \tag{2.58}$$

The local shunt-induced current perturbation was then evaluated as

$$\epsilon_I = \frac{R_S}{R_{\text{int,loc}} + R_{L,\text{eq}} + R_S}. \tag{2.59}$$

The estimated current that would flow without the low-side shunt is

$$I_0 = \frac{I_{\text{meas}}}{1-\epsilon_I} = I_{\text{meas}} \frac{R_{\text{int,loc}} + R_{L,\text{eq}} + R_S}{R_{\text{int,loc}} + R_{L,\text{eq}}}. \tag{2.60}$$

Assuming that the equivalent external load remains unchanged over this small perturbation, the corrected load voltage is

$$V_{L,0} = I_0 R_{L,\text{eq}}. \tag{2.61}$$

Since the shunt is absent in the corrected estimate, this corrected load voltage is also the corrected PMFC terminal voltage. The corrected PMFC-side electrical power is therefore

$$P_{\text{PMFC,corr}} = I_0 V_{L,0} = I_0^2 R_{L,\text{eq}}. \tag{2.62}$$

Accordingly, pre-correction PMFC-side power refers to $P_{\text{PMFC,inst}}$ from Eq. 2.52, i.e., the PMFC-side electrical power in the instrumented circuit including the shunt drop. Post-correction PMFC-side power refers to $P_{\text{PMFC,corr}}$ from Eq. 2.62, i.e., the estimated PMFC-side electrical power that would have been obtained without the 6.7 Ω low-side shunt. This correction is applied to the shunted intermittent-cell power measurements used in the long-term analysis; continuously loaded reference measurements that were not acquired through the same shunted current-sensing path are left unchanged.

Power consumption of the intermittent harvesting interface

As a final step of the intermittent harvesting platform characterization, the power consumption of the interface was measured. The circuit was externally supplied between 2.1 V and 4.5 V, a range selected from the operating requirements of the bipolar supply generating the ±2 V rails. For each hardware configuration, the supply voltage and current were recorded with a Keithley 2470 Source Measure Unit. The reported value at each supply voltage is the average over a stable time window after discarding the initial transient. The characterization was repeated for the BQ25505-only interface, the BQ25505 with bipolar supply, the BQ25505 with bipolar supply and intermittent timer, and the same interface with the voltage- and/or current-sensing front ends enabled.

Fig. 2.12 compares the proposed intermittent interface with two BQ25505-based reference configurations: the BQ25505 alone

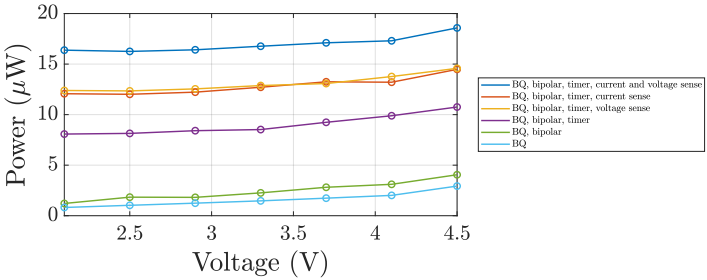


Figure 2.12: Power consumption of the reference BQ25505-based PMFC harvesting interface and of the proposed intermittent harvesting interface as a function of supply voltage. The proposed interface builds on the BQ25505 stage by adding the bipolar supply, intermittent timer, and optional voltage- and current-sensing front ends, which can be disabled to reduce power consumption when monitoring is not required.

and the BQ25505 with the bipolar supply. The BQ25505-only interface consumes approximately 0.8–2.9 μW over the tested 2.1–4.5 V supply range, while adding the bipolar supply raises the consumption to approximately 1.2–4.0 μW. These configurations represent the baseline consumption of a converter-based PMFC harvesting interface, with or without auxiliary bipolar rails required by analog front ends. The proposed intermittent extension, consisting of the BQ25505, bipolar supply, and timer, consumes approximately 8.0–10.8 μW. Thus, controlled intermittent operation adds about 6.8–7.9 μW with respect to the BQ25505-only interface, or about 6.5–6.8 μW with respect to the BQ25505-plus-bipolar baseline. Enabling the proposed voltage- and current-sensing front ends increases the consumption to approximately 16.3–18.6 μW, corresponding to an additional 7.8–8.3 μW over the timer-based intermittent interface. These sensing front ends are required for power tracking and characterization, but can be disabled in autonomous operation to reduce the steady-state overhead. These values must be interpreted with respect to the harvested-power levels reported in Sec. 3.2. In the tested PMFCs, the average harvested power is in the order of several tens to more than one hundred microwatts. Therefore, the timer-based intermittent extension represents a non-negligible but still compatible overhead for cells operating in the higher-power range. For example, when the intermittently operated

cells reach average harvested powers of 67.6 μW and 122.4 μW, the 8.0–10.8 μW consumption of the timer-based interface corresponds to approximately 12%–16% and 7%–9% of the harvested power, respectively. Conversely, for weaker cells the same overhead becomes a larger fraction of the available power budget: at 19.4 μW, the complete timer-based interface consumes approximately 41%–56% of the harvested power, while the timer extension alone accounts for approximately 35%–41%.

The discrete prototype is therefore compatible with the characterized PMFC power range, but further optimization is required for very weak cells. In autonomous deployments, the voltage- and current-sensing front ends should be disabled or duty-cycled when continuous power tracking is not required. Future implementations could further reduce the always-on overhead by replacing the MAX40018-based Schmitt-trigger timer with nanopower timing circuitry, for example a TLV3691-class comparator, and by using high-value feedback networks, low-leakage switching, and more integrated power-management circuitry. In particular, VLSI integration of the timer, bipolar supply, and sensing front ends into a PMFC-specific power-management interface could reduce the overhead associated with discrete components and make autonomous intermittent harvesting sustainable even for weaker PMFCs.

2.3 Summary

This chapter established the electrochemical and system-level motivation for controlled intermittent PMFC harvesting and presented the corresponding harvesting interface. The proposed architecture provides explicit control of the extraction timing while remaining compatible with converter-based energy delivery and PMFC monitoring. The following chapter evaluates the platform experimentally, first in terms of harvested-power response and subsequently through electrode-resolved impedance measurements under different loading histories.

Chapter 3

Experimental Evaluation of Intermittent PMFC Harvesting

Parts of this chapter are adapted from the author's previously published works on Plant Microbial Fuel Cells and battery-free electrochemical monitoring [20, 9, 25, 26].

This chapter experimentally evaluates the controlled intermittent harvesting interface developed in Chapter 2 and its effect on PMFC operation. The objective is twofold. First, we assess whether the proposed converter-based intermittent harvesting interface can reproduce the power-enhancement trends reported in intermittent MFC and PMFC studies, while remaining compatible with battery-free PMFC nodes. Second, we investigate whether the imposed loading profile leaves measurable signatures in the impedance response of the cells and, in particular, whether different loading histories produce distinguishable electrochemical responses at the anode and cathode.

The evaluation therefore combines power measurements and EIS measurements. Power measurements quantify the energy extracted under continuous and intermittent loading, and are used to iden-

tify favorable intermittent timing conditions. EIS is then used as a complementary, non-destructive observation method to compare the impedance response of PMFCs operated under intermittent harvesting, continuous harvesting, and unloaded conditions. The goal of the EIS analysis is not to identify the biological origin of the observed power variations, but to determine whether different loading histories are reflected in the spectra and in the fitted impedance parameters.

The experimental protocol was designed to compare continuous and intermittent PMFC loading while keeping the main electrical operating conditions fixed. The loaded cells were operated at a fixed cell voltage selected from preliminary power characterization, whereas the EIS measurements were performed under polarization conditions representative of the same harvesting operation.

The experimental evaluation addresses these two aspects in turn. First, using this platform, we observe higher PMFC power production under selected intermittent-harvesting timing conditions. Preliminary frequency exploration at 50% duty cycle identifies seconds-scale periods, especially around 5–10 s, as favorable in the tested cells. In a 9-day experiment, the intermittently operated cells were first driven with a 20 s period for six days and then with a 10 s period for the final three days. After this transition, Cell I reached an average corrected PMFC-side power of 122.4 μW, corresponding to a 159% increase with respect to its 20 s phase and a 185% increase with respect to the continuously loaded reference cell over the same interval. Cell H increased from 19.4 μW to 67.6 μW, corresponding to a 248% increase, and produced 57% more power than the continuous reference during the 10 s phase. These results suggest that benefits in absolute power output, as well as improvements in initially weaker cells, can be obtained in a practical two-electrode harvesting architecture, consistently with trends reported in intermittent-polarization and intermittent-loading studies.

The second part of the evaluation uses impedance-based state observation to evaluate the electrochemical consequences of different PMFC loading profiles. Building on the battery-free EIS feasibility demonstrated in [20], we compare PMFCs operated under intermittent harvesting, continuous harvesting, and unloaded conditions, and show that the imposed loading profile produces measurable changes in the EIS response. This supports the use of EIS as a non-destructive diagnostic method for PMFC-powered nodes and

motivates its future use in PMFC-specific MPPT and biosensing strategies.

The chapter does not claim to fully identify the biological origin of the observed impedance changes. Rather, it establishes a battery-free-compatible platform for applying controlled intermittent harvesting and observing its electrochemical consequences. The power results show substantially higher harvested power under selected intermittent timing conditions, while the EIS results indicate that loading history is reflected in the PMFC impedance response. Complementary microbiological, electrochemical, and plant-physiology measurements are therefore required to determine whether these changes arise from biofilm restructuring, microbial activity, transport limitations, hydration, cathode behavior, or plant-related effects. This motivates future work in which harvesting, stabilization, and state observation are treated as coupled components of the same PMFC energy-management problem.

The chapter is organized as follows. Sec. 3.1 describes the PMFCs, the loading profiles, the power-acquisition setup, and the EIS measurement protocol. Sec. 3.2 reports the harvested-power results, starting from a preliminary timing exploration and then moving to a 9-day comparison between intermittent and continuous loading. Sec. 3.3 analyzes the impedance spectra acquired during the same experiment, focusing on loading- and hydration-related trends rather than on a complete mechanistic interpretation. Finally, Sec. 3.4 summarizes the main findings of the experimental evaluation and discusses the remaining limitations and future directions for PMFC-specific energy management.

3.1 Experimental setup and measurement protocol

3.1.1 PMFC construction and start-up

Four PMFCs were assembled using the same nominal geometry, materials, and start-up protocol. Each cell used rectangular carbon-felt electrodes (9 x 14 cm) and plastic containers without drainage holes (19 cm diameter, 17 cm height). The anode was buried at a depth of approximately 10 cm, while the cathode was placed on the soil surface, as shown in Fig. 3.2. The lower half of each container, sur-

rounding the anode, was filled with a 50:50 mixture of commercial soil and soil collected from the university campus, the latter being used as bacterial inoculum. The upper half of the cell was filled with commercial soil only.

To start up the PMFCs, the cells were flooded with water and left under open-circuit conditions until the open-circuit voltage reached approximately 600 mV. After start-up, water was replenished whenever the top soil layer started to visibly lose moisture. Before the loading experiment, all cells were kept unloaded for seven days in order to promote comparable initial conditions. These controls reduce construction-related differences, but they do not eliminate biological and environmental variability, including local electrode–soil contact, microbial colonization, oxygen gradients, hydration distribution and substrate availability.

3.1.2 Loading protocols and power acquisition

After the initial conditioning phase, the four cells were assigned to three loading profiles. Cells I and H were operated intermittently using the intermittent-harvesting architecture presented in Sec. 2.2, Cell C was operated continuously, and Cell U was kept unloaded as a control. For the intermittent cells, the duty cycle was fixed at 50%, while the switching period was varied according to the experimental phase. In the preliminary timing exploration, periods between 5 s and 1000 s were tested to identify a suitable operating range. In the long-term experiment of Sec. 3.2.3, the intermittent cells were operated with a 20 s period during the first six days and with a 10 s period during the final three days. A complete joint optimization of duty cycle and period is left as future work.

The continuous and intermittent cells were loaded at a fixed cell operating voltage of 270 mV. This value was selected because preliminary characterization placed it close to the maximum-power-point region. Active MPP tracking was deliberately avoided, since it would have introduced an additional time-varying control variable and would have made the comparison between continuous and intermittent loading less direct. For the continuously operated cell, the 270 mV operating point was imposed using a Keithley 2450 source-measure unit. For the intermittently operated cells, the same cell-voltage target was imposed during the ON phase through the input-modulation function of the BQ25505, using the external ref-

erence provided by a Keithley 2470 source-measure unit.

Power measurements on the intermittently operated cells were acquired using the setup shown in Fig. 3.1. A custom LabVIEW application running on a portable computer controlled the acquisition of the cell voltages and of the current delivered by the active cell. The measurements were performed through the voltage- and current-sensing front ends described in Sec. 2.2.2 and acquired using a National Instruments NI USB-6259 data logger. Cell voltages were monitored with unity-gain buffer stages, while the active-cell current was measured with a low-side shunt followed by a two-op-amp instrumentation-amplifier stage, whose noise analysis and calibration are reported in Sec. 2.2.2. Since this low-side shunt is inserted in the PMFC return path, the acquired current corresponds to the instrumented current I_{meas}, rather than to the ideal uninstrumented current.

For the long-term power analysis, the active-sample current, load voltage, timestamps and cell identifier were retained for post-processing. These quantities were used together with the EIS-derived local source-resistance estimates to compute the PMFC-side electrical power in the instrumented configuration, $P_{\mathrm{PMFC,inst}}$, and to estimate the corresponding shunt-corrected PMFC-side power, $P_{\mathrm{PMFC,corr}}$, that would have been obtained without the 6.7 Ω low-side shunt.

During this experimental campaign, the intermittent interface, which is normally intended for self-powered operation, was supplied externally at 3 V using a GW Instek GPS-1850D power supply to ensure uninterrupted operation of the control and acquisition circuitry. The consumption of the interface itself is characterized separately in Sec. 2.2.2.

3.1.3 EIS measurement protocol

EIS measurements were designed to probe the electrochemical response of the PMFCs under conditions representative of energy-harvesting operation. During harvesting, the power-management interface imposes a cell voltage and therefore polarizes both electrodes; since the impedance spectrum depends on the DC polarization point, spectra acquired at open circuit may not describe the response experienced during energy extraction. This dependence was recently highlighted by Koo *et al.*, who showed that anodic impedance parameters vary with the imposed anode potential in a

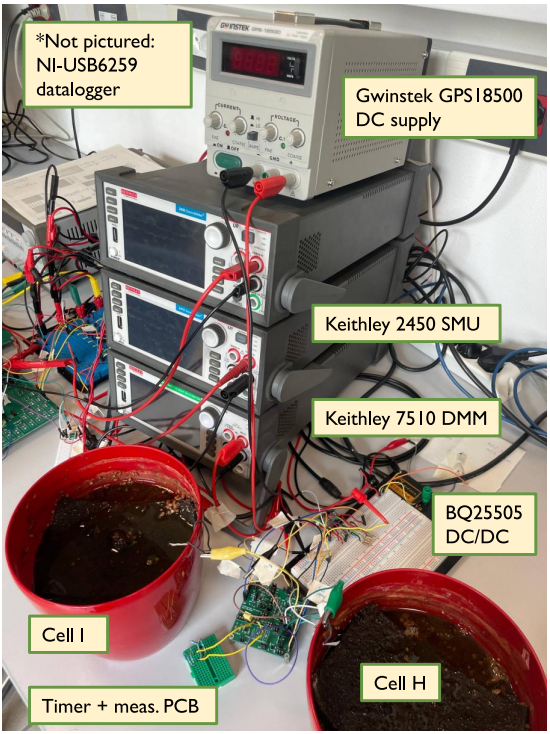


Figure 3.1: Intermittent harvesting setup used for Cells I and H.

continuous-flow MFC [78]. For this reason, the anode and cathode potentials reached during two-electrode harvesting at the imposed 270 mV cell voltage were used as reference conditions for the corresponding three-electrode EIS measurements. This provides an experimental link between the two configurations: the harvester establishes the coupled full-cell operating condition, while EIS subsequently probes the anode and cathode separately at polarization points derived from that condition. The EIS measurement therefore does not reproduce two-electrode harvesting directly; instead, the electrode potentials reached during harvesting are used as the polarization conditions for separate anode and cathode characterization.

EIS measurements were performed using a PalmSens4 electrochemical interface, as shown in Fig. 3.2. A BASi MF-2056 Ag/AgCl reference electrode filled with 3 M KCl was used as reference. The anode and cathode are denoted by A and C, respectively, and each EIS measurement is identified by cell name, electrode, and loading day. The first measurement was acquired before the loading test started and is therefore referred to as Day 0.

For the loaded cells, the anode and cathode potentials were sampled during operation and then used as the DC polarization points for the potentiostatic EIS measurements. For the unloaded control cell, the electrode polarization points were obtained after temporarily loading the cell at 270 mV for 5 min. In this way, the spectra were acquired under conditions representative of the harvesting state, rather than only under open-circuit relaxation.

Before each impedance scan, a 30 s equilibration time was introduced to reduce the effect of initial capacitive transients. The EIS sweep was then performed from 50 mHz to 1 MHz, using 61 frequency points and a 10 mV AC perturbation amplitude. The frequency range was selected from a battery-free perspective: as shown in the previous battery-free EIS implementation [20], impedance monitoring can be compatible with PMFC-level energy production when sweep duration and energy cost are bounded. The lower frequency bound was therefore kept in the tens-of-millihertz range, preserving access to slow impedance features while avoiding the excessive acquisition time and energy cost associated with lower-frequency measurements. A differential PicoScope 4000 was used in parallel to monitor the electrode voltage evolution during the tests.

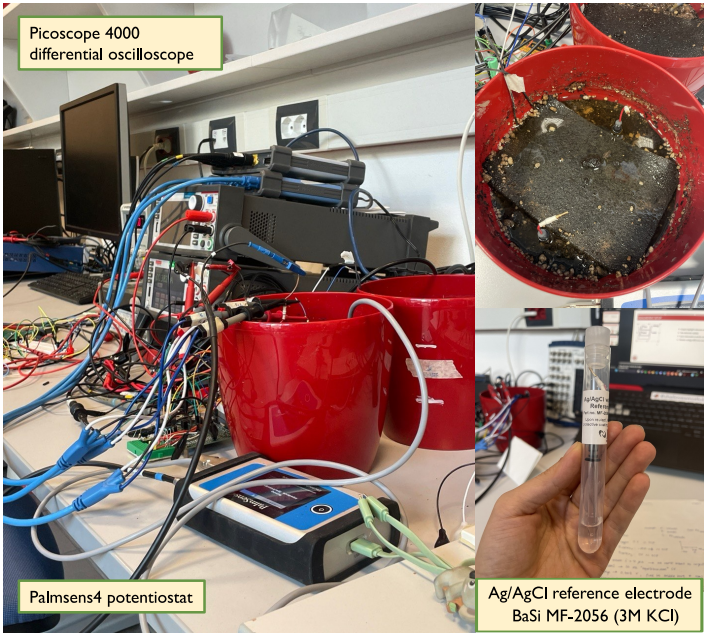


Figure 3.2: EIS measurement setup used in parallel with the intermittent harvesting setup shown in Fig. 3.1.

3.1.4 Measurement timeline and hydration events

EIS measurements were acquired at the beginning of the experiment and then every 24 hours for the first three days. After Day 3, the next EIS measurement was performed only after three additional days, in order to observe cell evolution while reducing the disturbance introduced by repeated EIS measurements.

The timing of hydration events was explicitly considered in the analysis. Water was added to all cells on Day 0, two hours before the first EIS measurement, and again after the Day 3 measurement. Measurements labelled Day 3 and Day 3W therefore correspond to the states immediately before and immediately after wetting on the third day, respectively, when both spectra were available.

For the intermittently operated cells, additional watering was required between Day 1 and Day 2 because these cells appeared to dry out faster, with Cell I showing the driest condition. The transition from Day 1 to Day 2 was therefore used to evaluate the response following the visibly dry Day 1 state. A larger water volume was added on Day 3 to reduce uneven drying across the cells. Since water was added to all cells after the Day 3 measurement and no additional EIS measurement was performed before Day 6, the Day 6 spectra must be interpreted as the combined result of rehydration and continued, unperturbed operation over the following three days.

3.2 Power response to intermittent harvesting

3.2.1 Shunt-burden correction of PMFC-side power measurements

The shunt-burden correction was applied to the intermittently loaded cells to estimate the PMFC-side electrical power that would have been obtained without the 6.7 Ω low-side current-sensing shunt. The correction outcome is quantified for the 9-day experiment, since this is the most complete long-term dataset and includes the segment-wise current, voltage, and EIS-derived source-resistance estimates required by the correction procedure.

Fig. 3.3 compares the instrumented PMFC-side power, $P_{\mathrm{PMFC,inst}}$, with the corrected estimate, $P_{\mathrm{PMFC,corr}}$, for Cells I and H. Points

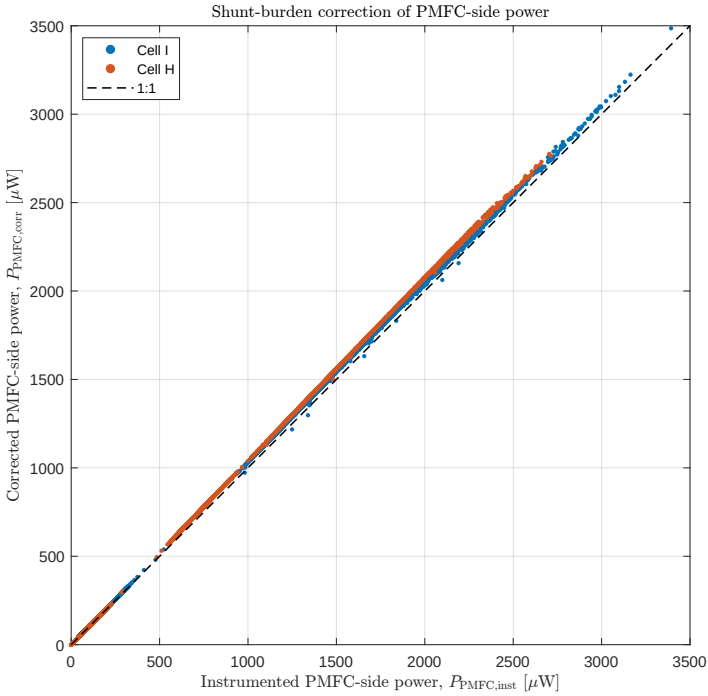


Figure 3.3: Effect of the shunt-burden correction on the PMFC-side electrical power of the intermittently loaded cells during the 9-day experiment.

above the 1:1 line indicate the increase introduced by the correction. The separation from the 1:1 line becomes more visible at higher power, where the measured current is larger and therefore both the shunt voltage drop, $I_{\mathrm{meas}}R_S$, and the shunt-related power term, $I_{\mathrm{meas}}^2 R_S$, increase. Tab. 3.1 reports the corresponding current perturbation, ϵ_I, defined in Eq. 2.59, and the resulting PMFC-side power increase.

Across the 9-day experiment, the mean current perturbation ϵ_I was 1.34% for Cell I and 0.68% for Cell H. Since $\epsilon_I = (I_0 - I_{\mathrm{meas}})/I_0$, this means that the measured current was on average 1.34% and 0.68% lower than the estimated no-shunt current, respectively. The corresponding mean increases in PMFC-side power were 1.23% and 0.66%, respectively. The correction therefore remains small and does not change the qualitative outcome of the experiment, although

Table 3.1: Active-sample statistics of the shunt-burden correction for the intermittently loaded cells during the 9-day experiment.

Cell	Current perturb., ϵ_I [%]		Power increase [%]	
	Mean	Max	Mean	Max
I	1.34	10.08	1.23	3.15
H	0.68	10.21	0.66	4.01

local maxima occur during higher-current operating points.

Unless explicitly stated otherwise, the PMFC-side power values reported for the intermittently loaded cells in Sec. 3.2.3 correspond to the corrected estimates $P_{\mathrm{PMFC,corr}}$. The instrumented values $P_{\mathrm{PMFC,inst}}$ are retained in Fig. 3.3, Tab. 3.1, the pre-correction columns of Tab. 3.2, and the complementary time-domain comparison in Fig. 3.7 to document the magnitude of the correction.

3.2.2 Preliminary timing exploration

Before starting the long-term intermittent-harvesting experiment, a short preliminary exploration was performed to identify a suitable switching period. The duty cycle was fixed at 50%, while the period was varied. In the first test, three periods spanning two decades were compared over a 6000 s interval: 10 s, 100 s, and 1000 s. These periods were obtained using $R_{F1} = R_{F2} = 82$ MΩ as timer feedback resistors and timing capacitances of 22 nF, 220 nF, and 2200 nF, respectively.

The corresponding average harvested powers are reported in Fig. 3.4. The 10 s period produced the highest average power, 83 μW, followed by 64 μW for the 100 s period and 51 μW for the 1000 s period. Therefore, decreasing the period from 1000 s to 100 s increased the average harvested power by approximately 25%, while decreasing it further to 10 s produced an increase of approximately 63% with respect to the 1000 s case and 30% with respect to the 100 s case. This first exploration therefore indicated that shorter periods were more favorable for average power production.

A second preliminary test was then performed over a longer duration of approximately 12 h, focusing on periods close to 10 s. In this test, 5 s, 10 s, and 20 s periods were compared. These periods were obtained using the same 82 MΩ timer feedback resistors and timing capacitances derived from 22 nF: two 22 nF capacitors

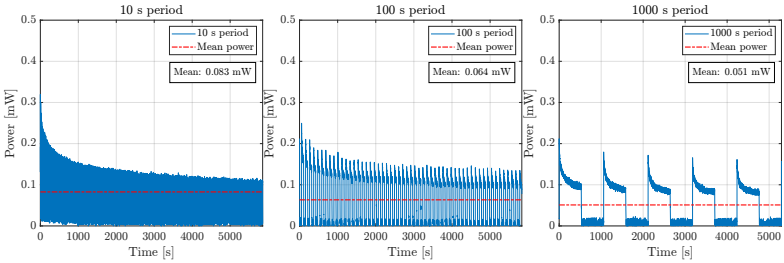


Figure 3.4: Preliminary intermittent-harvesting frequency exploration over 6000 s, comparing 10 s, 100 s, and 1000 s periods at 50% duty cycle.

in series for 5 s, a single 22 nF capacitor for 10 s, and two 22 nF capacitors in parallel for 20 s.

The results, shown in Fig. 3.5, indicate that the 10 s and 5 s periods produced very similar average powers, 37 μW and 36 μW, respectively. The 20 s period produced a lower average power of 28 μW. Thus, the 10 s period increased the average harvested power by approximately 32% with respect to the 20 s case, while the 5 s period produced an increase of approximately 29%. The difference between 5 s and 10 s was small, with the 10 s period producing only about 3% more average power.

The average values obtained in the 12 h test are lower than those obtained in the shorter 6000 s exploration because the longer acquisition includes the full relaxation of the initial transient. In particular, the power traces show a large capacitive transient at the beginning of the measurement, followed by gradual settling toward a plateau after approximately 4 h. Therefore, the 12 h averages are more representative of regime power production, whereas the shorter test is more strongly influenced by the initial transient response.

These preliminary tests are still limited in duration and cannot capture long-term biological adaptation effects. Nevertheless, they indicate that the most promising switching periods for increased average harvested energy lie around 5–10 s, with 20 s providing a slightly less favorable but nearby operating point.

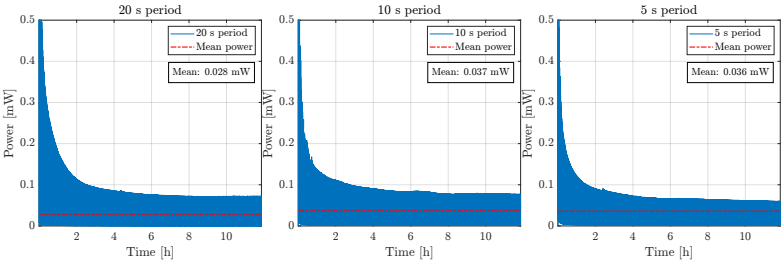


Figure 3.5: Extended preliminary frequency exploration over approximately 12 h, comparing 5 s, 10 s, and 20 s periods at 50% duty cycle.

3.2.3 Nine-day continuous and intermittent loading experiment

Based on these results, a longer 9-day test was performed using newly started Cells I and H, operated intermittently, and Cell C, operated continuously. The aim was to compare intermittent harvesting with continuous loading over several days and to evaluate whether moving from a slightly sub-optimal period to the most promising period identified in the preliminary tests would produce a measurable increase in average intermittent power. A separate continuously loaded control cell was used under the assumption that intermittent operation may induce long-term biofilm conditioning, as reported for periodically polarized electroactive biofilms [10]. Using the same cell as both continuous and intermittent control could therefore confound the comparison, since the later operating phase might be affected by previous biofilm adaptation. During the first six days, the intermittent cells were operated with a 20 s period. During the last three days, the period was reduced to 10 s. EIS measurements were also collected during the first part of the experiment, as discussed in Sec. 3.1.

Fig. 3.6 shows the time evolution of the harvested power during the 9-day loading experiment. For the intermittently loaded cells, the plotted values correspond to the corrected PMFC-side power, $P_{\mathrm{PMFC,corr}}$, as defined in the shunt-burden correction analysis at the beginning of Sec. 3.2. As shown in Fig. 3.7, the correction remains small over the 9-day experiment and does not change the qualitative trends discussed below. The numerical comparison below therefore refers to the corrected PMFC-side values reported in Tab. 3.2.

Table 3.2: Average PMFC-side electrical power during the 9-day experiment before and after shunt-burden correction. The pre-correction values correspond to the instrumented PMFC-side power $P_{\mathrm{PMFC,inst}}$, whereas the post-correction values correspond to the corrected estimate $P_{\mathrm{PMFC,corr}}$. The intermittent period changed from 20 s to 10 s on Day 7, while Cell C remained continuously loaded.

Cell	Loading	Days 1–6		Days 7–9		Increase	Days 7–9 vs. C
		Pre-corr. (μW)	Post-corr. (μW)	Pre-corr. (μW)	Post-corr. (μW)	Post-corr.	Post-corr.
I	Intermittent 20 → 10 s	46.7	47.2	120.1	122.4	+159%	+185%
H	Intermittent 20 → 10 s	19.3	19.4	66.7	67.6	+248%	+57%
C	Continuous	32.8	32.8	43.0	43.0	+31%	Reference

During the first six days, the 20 s intermittent period produced 47.2 μW for Cell I and 19.4 μW for Cell H, compared with 32.8 μW for the continuously loaded Cell C. Cell H therefore produced about 59% less power than Cell I during this first phase, highlighting the strong cell-to-cell variability of PMFCs despite identical nominal assembly, start-up procedure, and loading voltage. In addition, although the cells were watered following the same general procedure, the intermittently operated cells appeared to dry faster. This caused a visible power decrease in Cell I at approximately 40 h of loading time, between the Day 2 and Day 3 EIS measurements.

After the transition from 20 s to 10 s during the last three days, both intermittently operated cells showed a marked improvement in PMFC-side power. Cell I reached 122.4 μW, corresponding to a 159% increase with respect to its 20 s phase, while Cell H increased to 67.6 μW, corresponding to a 248% increase. The continuously loaded Cell C also increased during the same interval, from 32.8 μW to 43.0 μW, despite undergoing no configuration change. Therefore, the improvement observed in the intermittent cells cannot be attributed solely to the period change, since biofilm stabilization, hydration dynamics, ambient conditions, or other biological effects may also have contributed. Nevertheless, the growth rates of the intermittent cells were much larger than that of Cell C, which increased by 31%, and during the final interval Cells I and H exceeded Cell C by 185% and 57%, respectively.

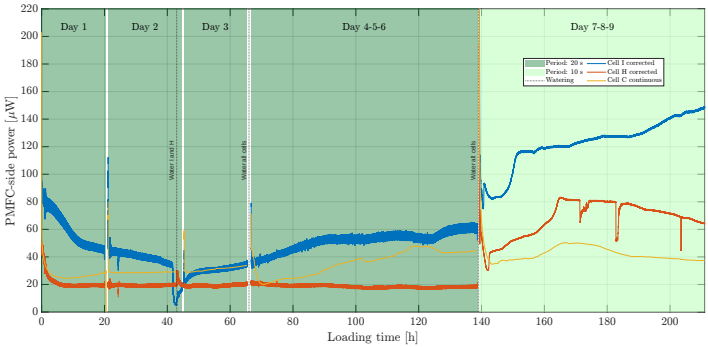


Figure 3.6: Power evolution during the 9-day loading experiment. Cells I and H were operated intermittently, while Cell C was operated continuously. The intermittent period was 20 s during the first six days and 10 s during the last three days.

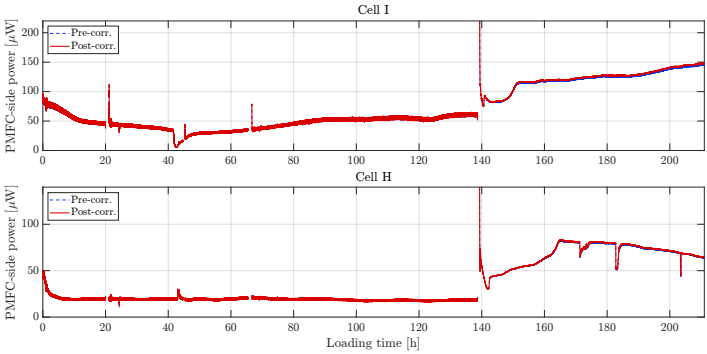


Figure 3.7: Effect of the shunt-burden correction on the PMFC-side electrical power of the intermittently loaded cells during the 9-day experiment. Dashed blue traces show the instrumented PMFC-side power, $P_{\mathrm{PMFC,inst}}$, obtained with the 6.7 Ω low-side shunt present. Solid red traces show the corrected estimate, $P_{\mathrm{PMFC,corr}}$, corresponding to the same equivalent external load after removing the shunt effect. The close overlap between pre- and post-correction traces indicates that the correction is small and does not change the observed power trends.

3.2.4 Discussion of power trends

Overall, these results support the hypothesis that controlled intermittent extraction can improve PMFC energy harvesting when the timing is properly selected, and that power-enhancement trends reported for intermittent MFC operation can be reproduced with a battery-free-compatible two-electrode harvesting architecture. In the present experiment, the most favorable preliminary condition was obtained with a 10 s period and 50% duty cycle, which led both intermittently operated cells to outperform the continuously loaded reference during the final phase. At the same time, the increase observed in the continuous reference confirms that biological and environmental dynamics remain important confounding factors. The present data therefore support controlled intermittency as a promising strategy, but do not allow the contribution of the 20 s to 10 s period change to be quantified separately from cell variability, hydration, or biofilm evolution. The proposed interface remains useful both as a harvesting architecture and as an experimental platform, because it imposes period and duty cycle as explicit circuit-level variables.

Previous studies also reported measurable gains over constant or steady-state loading when the switching parameters were properly selected, although generally at different timescales and outside a battery-free harvesting setup. Coronado et al. reported a 22–43% increase in MFC power output using pulse-width-modulated external resistance at switching frequencies above 100 Hz. Their duty-cycle optimization was performed at 500 Hz, corresponding to a 2 ms period, with maximum power obtained around 95% duty cycle [63]. Liang et al. reported an 18.4% increase in maximum power density and a 32.7% increase in long-term average power density using transient-state regulation compared with steady-state operation. Their best condition was obtained at 1 Hz, corresponding to a 1 s period, and 75% duty cycle [62]. The condition identified here is therefore slower than those reported in these power-oriented studies, but remains consistent with the broader conclusion that intermittent operation can improve MFC power production when the period and duty cycle are properly selected.

The 5–10 s timescale observed in this work is also compatible with studies emphasizing short recovery intervals. Grondin et al. showed that seconds-scale intermittent connection of the external

load can approach the power obtained under optimal continuous loading [13], while Gardel et al. reported that duty cycling affects cumulative charge and current generation in multi-anode environmental MFCs using switching intervals on the order of seconds to tens of seconds [12]. Moreover, the use of a 50% duty cycle agrees qualitatively with the long-term electroactive-biofilm study of Zhang et al., where lower-duty-cycle periodic polarization improved faradaic current production and total charge recovery over several weeks [10]. However, similar timing does not imply equivalent electrochemical conditions, since the present work controls the full-cell voltage in a two-electrode harvesting configuration rather than an individual electrode potential. The agreement in duty cycle and characteristic timescale is therefore encouraging, but does not demonstrate that the optimum identified under three-electrode polarization transfers directly to two-electrode harvesting.

These results should therefore be considered preliminary but encouraging. Although a complete optimization of period and duty cycle was outside the scope of this preliminary test, the developed architecture enables such exploration directly in an embedded and self-powered harvesting setup. Since the switching period and duty cycle are configurable at the circuit level, the platform can be used to identify suitable intermittent-harvesting conditions under realistic PMFC operation, outside of a laboratory environment. Once an operating point is selected, the same architecture is ready to be deployed in real-world battery-free applications powered by PMFCs, without requiring a separate electrochemical control setup.

3.3 Impedance response to loading history

The power measurements discussed in Sec. 3.2 quantify how much energy is extracted under each loading profile, but they do not describe how the PMFC impedance response evolves under those same conditions. For this reason, EIS is used here as a complementary, non-destructive observation method. The objective is not to identify the biological or electrochemical origin of the observed power variations, nor to derive a complete mechanistic model of the PMFCs. Rather, the goal is to assess whether intermittent harvesting, continuous harvesting, and unloaded operation are associated with

measurable differences in the impedance spectra. This comparison therefore tests whether a controlled temporal loading profile applied at the full-cell level leaves measurable electrode-level signatures, as would be expected from the loading-history dependence reported in controlled electrochemical studies.

In this sense, the EIS study provides insight into the electrochemical signatures associated with different loading histories, rather than a causal explanation of the power trends. We therefore compare the spectra of PMFCs operated under the three loading profiles and analyze how their impedance response evolves over the multi-day experiment. Particular attention is given to hydration events, since soil moisture strongly affects ionic conduction and transport processes and may therefore introduce impedance changes that are not directly attributable to the electrical loading profile.

The analysis proceeds in three steps. First, the impedance spectra are interpreted using a common equivalent-circuit model for both anode and cathode, allowing fitted parameters to be compared across cells, electrodes, and operating conditions. Second, the effect of electrical loading is evaluated by comparing continuously loaded, intermittently loaded, and unloaded PMFCs. Third, the effect of hydration is examined using spectra acquired before and after watering events. Throughout the analysis, the fitted parameters are interpreted comparatively rather than as unique physical descriptors. Determining whether the observed changes arise from biofilm restructuring, microbial activity, hydration, transport limitations, cathode behavior, or plant-related effects requires complementary validation measurements beyond the scope of this chapter.

3.3.1 Equivalent-circuit model for impedance interpretation

The same equivalent-circuit topology was used for both the anode and cathode in order to compare the two electrodes through a consistent impedance decomposition. The selected circuit consisted of a series ohmic resistance, one high-frequency R–CPE branch, and a second mixed CPE/finite-diffusion branch. Following standard EIS practice, this topology was used to separate the fitted impedance into ohmic, distributed interfacial charge-transfer/double-layer, and concentration- or diffusion-related contributions [79, 80, 48]. For MFC electrodes, analogous combinations of series resistance, R–

CPE branches, and Warburg-type elements are commonly used to distinguish electrolyte/electrode ohmic contributions, biofilm or cathode interfacial arcs, and low-frequency transport limitations [48, 81, 82, 83]. The circuit was therefore intended to provide a consistent and physically interpretable decomposition of the main impedance contributions, rather than a unique description of the underlying electrode processes. The fitted impedance can be written in simplified form as:

$$Z(\omega) = R_s + (R_1 \parallel Z_{\mathrm{CPE},1}) + [Z_{\mathrm{CPE},2} \parallel (R_2 + Z_{W,o})], \tag{3.1}$$

where the impedance of a constant phase element is

$$Z_{\mathrm{CPE}} = \frac{1}{Q_0(j\omega)^n}, \tag{3.2}$$

and the finite-length open Warburg element is expressed as

$$Z_{W,o} = R_W \frac{\coth\left(\sqrt{j\omega\tau_D}\right)}{\sqrt{j\omega\tau_D}}. \tag{3.3}$$

The series resistance R_s represents the ohmic contribution of the electrolyte, electrode/electrolyte contacts, current collectors, and separator when present. The first branch, $R_1 \parallel Z_{\mathrm{CPE},1}$, represents the main interfacial process. At the anode, this branch was assigned to the biofilm/electrode interface, including double-layer charging and extracellular electron-transfer processes. At the cathode, it was assigned to the dominant oxygen-reduction/interfacial arc. In an ideal electrode, this contribution would correspond to $R_{charge\ transfer} \parallel C_{double\ layer}$, where the resistance describes the kinetic difficulty of transferring electrons to dissolved oxygen at the electrode/electrolyte interface and the capacitance describes double-layer charging. At the cathode, a larger R_1 indicates a larger apparent ORR/interfacial limitation, which may result from slower oxygen-reduction kinetics, reduced oxygen availability, or poorer ionic contact at the electrode/electrolyte interface. At the anode, an increase in R_1 indicates a larger apparent biofilm/electrode interfacial resistance, which may arise from less efficient extracellular electron transfer, poorer biofilm/electrode contact, reduced local substrate availability, or changes in biofilm conductivity [84].

The interfacial storage contribution was represented by a CPE - and not by an ideal capacitor - because both electrodes are porous,

heterogeneous, and spatially non-uniform [84]. The biofilm-covered anode may contain variations in biofilm thickness, active biomass, hydration, ionic conductivity, and electron-transfer pathways, while the oxygen-reducing cathode may show non-uniform wetting, oxygen gradients, and spatially distributed reaction sites. These features generate a distribution of local time constants rather than a single ideal capacitive response; therefore, $Z_{\mathrm{CPE},1}$ was used to describe the non-ideal double-layer behavior, as commonly done for MFC electrodes [48].

The CPE coefficient Q_1 controls the magnitude of the capacitive-like interfacial response. A larger Q_1 indicates greater apparent charge-storage capacity. At the anode, this may reflect a larger electrochemically active biofilm/electrode interface, stronger double-layer or pseudo-capacitive storage, and biological capacitance associated with active electricigenic biomass [85]. The exponent n_1 describes the degree of non-ideality of the interfacial response. Values close to 1 indicate a more ideal and uniform capacitive behavior, whereas lower values indicate a more distributed or heterogeneous interface. At the anode, a lower n_1 may reflect heterogeneous biofilm structure, non-uniform hydration, spatially distributed electron-transfer sites, or variations in ionic transport through the biofilm. At the cathode, it may reflect uneven wetting, non-uniform oxygen access, or spatial variation in ionic resistance. Thus, R_1, Q_1, and n_1 were interpreted together: R_1 describes the apparent kinetic/interfacial resistance, Q_1 the magnitude of interfacial storage, and n_1 the degree of distribution or heterogeneity of the process.

The second branch was interpreted as a slower distributed interfacial/transport process. At the anode, this contribution may include slower biofilm charge-transfer processes, substrate transport, proton or buffer transport, and concentration gradients within the hydrated biofilm/electrode region. At the cathode, it may include slower oxygen-reduction pathways, oxygen and water transport and proton or ionic transport. Therefore, this branch was not assigned to a single elementary process, but to the combined low-frequency response of interfacial kinetics and mass transport [84]. Within this branch, R_2 represents an apparent resistance associated with slower processes that are not resolved by the first R–CPE arc. At the anode, larger R_2 may indicate slower biofilm-mediated charge transfer, reduced substrate availability, poorer proton or buffer transport, or

stronger concentration gradients. At the cathode, larger R_2 may indicate slower ORR-related pathways, reduced oxygen access, wetting effects, or increased ionic resistance within the cathode layer. The parameters Q_2 and n_2 describe the magnitude and dispersion of the slower distributed storage contribution. Unlike Q_1, which is associated with the main interfacial arc, Q_2 reflects low-frequency storage coupled to slower transport or reaction-distribution effects. At the anode, larger Q_2 may indicate greater participation of the hydrated biofilm volume or stronger distributed ionic/redox storage [85, 86]. At the cathode, it may indicate greater participation of wetted porous regions or catalyst/electrolyte domains. The exponent n_2 indicates how distributed this slow process is. Values approaching $n_2 \approx 0.5$ are particularly important because they indicate Warburg-like frequency dependence. In those cases, part of the diffusion-like response may be absorbed by $Z_{\mathrm{CPE},2}$, making the explicit Warburg parameters less identifiable. Therefore, R_2, Q_2, n_2, R_W, and τ_D were interpreted together rather than independently. Additionally, for the first and second R-CPE branches, characteristic frequencies f_{c1} and f_{c2} were calculated to compare the apparent timescales of the two interfacial processes:

$$f_c = \frac{1}{2\pi}\left(\frac{1}{RQ}\right)^{1/n} \tag{3.4}$$

The finite-open Warburg element was used to describe bounded diffusion-like transport. Other diffusion elements were tested (Gerischer, finite-short, R_2 and CPE_2 only), but only the finite-open Warburg yielded acceptable data fitting. The finite-open Warburg element corresponds to a finite transport domain with a blocking boundary, beyond which the transported species is not present. For the anode, this was interpreted as apparent transport within a finite hydrated biofilm or porous-electrode region, where substrate, protons, buffer species, or redox-active species may be redistributed only within a confined domain. For the cathode, the finite-open Warburg element may be interpreted as an apparent bounded oxygen-transport response within the soil/cathode region. In this system, oxygen is supplied from the aerated side but may only penetrate to a finite depth in the soil before reaching an oxygen-depleted region. This lower oxygen-depleted region does not behave as an ideal well-mixed oxygen reservoir and may therefore limit the propagation of the concentration perturbation. In this sense, the fitted finite-open Warburg

element can represent a finite oxygen-accessible domain bounded by a poorly replenished or effectively oxygen-depleted region. However, this interpretation should be considered phenomenological, because an oxygen-depleted boundary is not necessarily a perfectly reflective boundary; depending on whether it behaves as a sink or as a non-replenishing barrier, finite-short or Gerischer-type descriptions may also be physically plausible. The Warburg resistance R_W describes the apparent amplitude of the bounded transport contribution. A larger R_W indicates a stronger low-frequency transport-related impedance, which may arise from more severe concentration gradients, poorer substrate or oxygen access, reduced hydration, or longer/more tortuous transport pathways. The time constant τ_D and the corresponding frequency $f_W = 1/(2\pi\tau_D)$ describe the apparent timescale of the bounded transport process: a larger τ_D therefore corresponds to a slower apparent transport process.

The low-frequency parameters were interpreted cautiously because R_2, $Z_{\mathrm{CPE},2}$, and $Z_{W,o}$ can become strongly coupled. Thus, when the low-frequency parameters approached fitting bounds or had uncertainties comparable to, or larger than, their fitted values, the corresponding f_{c2} or finite-Warburg contribution was considered non-identifiable and was not interpreted quantitatively.

Model fitting results are presented in Tabs. 3.3 and 3.4. The reliability of the fitted parameters was not uniform across the circuit. The most robust parameters were generally R_s, R_1, n_1, and f_{c1}, since they describe the high-frequency intercept and the main interfacial arc, which were well represented within the measured frequency range. In contrast, the low-frequency R-CPE and finite-diffusion branch parameters R_2, Q_2, n_2, R_W, and τ_D were less represented in the data, and thus less robust, because the spectra were acquired only down to 50 mHz, consistently with the lower-frequency bound identified for low-power EIS operation. These parameters were also sensitive to parameter coupling, especially when the slower R-CPE branch and the finite-diffusion contribution could not be clearly separated from the data. Second branch parameters were thus interpreted only if their standard error was lower than their estimated value, and are to be regarded as qualitative indicators, rather than quantitative.

Table 3.3: Finite open Warburg model fitting, anode.

ID	loading	HF_cut	Rs	R1	Q1	n1	R2	Q2	n2	Rw	tauD	f_c1	f_c2	f_W	RMS
unit		kHz	Ω	Ω	m $\frac{s^n}{\Omega}$	-	Ω	m $\frac{s^n}{\Omega}$	-	Ω	s	Hz	Hz	Hz	-
UA0	-	29.85	11.70 ± 0.04	1.39 ± 0.17	2.00 ± 0.64	0.63 ± 0.04	2.52 ± 1.34	299.59 ± 48.65	0.39 ± 0.05	7.03 ± 1.73	1.71 ± 0.77	1948.10	0.33	0.0931	0.04
UA1	-	23.44	8.43 ± 0.04	1.44 ± 0.12	2.70 ± 0.68	0.59 ± 0.03	2.19 ± 1.65	349.53 ± 47.48	0.44 ± 0.04	7.45 ± 1.92	1.72 ± 0.83	1964.77	0.29	0.0923	0.03
UA2	-	29.85	7.76 ± 0.02	1.32 ± 0.11	1.01 ± 0.24	0.71 ± 0.03	1.50 ± 0.62	313.77 ± 49.85	0.41 ± 0.05	6.30 ± 0.58	2.38 ± 0.56	1901.26	1.00	0.0668	0.03
UA3	-	8.91	3.35 ± 0.04	0.84 ± 0.07	1.98 ± 0.70	0.72 ± 0.05	1.37 ± 2.00	451.00 ± 66.83	0.51 ± 0.05	4.69 ± 2.91	1.00 ± 0.96	1201.45	0.41	0.1593	0.03
UA3W	-	11.35	4.54 ± 0.03	0.91 ± 0.10	1.32 ± 0.49	0.74 ± 0.05	1.62 ± 1.68	336.35 ± 57.05	0.45 ± 0.05	8.83 ± 1.16	2.25 ± 0.87	1394.87	0.61	0.0709	0.03
UA6	-	11.35	11.53 ± 0.09	1.64 ± 0.20	3.12 ± 1.01	0.54 ± 0.05	2.08 ± 0.98	329.80 ± 50.99	0.43 ± 0.05	6.08 ± 1.32	1.93 ± 0.74	2727.21	0.38	0.0823	0.03
HA0	int.	14.45	2.96 ± 0.04	2.29 ± 0.13	1.15 ± 0.24	0.69 ± 0.03	0.68 ± 1.92	203.86 ± 74.40	0.66 ± 0.09	6.15 ± 1.95	1.42 ± 0.96	913.50	3.14	0.1120	0.05
HA1	int.	14.45	1.99 ± 0.03	2.25 ± 0.06	0.90 ± 0.12	0.71 ± 0.02	0.32 ± 3.91	268.65 ± 132.39	0.79 ± 0.08	3.99 ± 4.86	0.62 ± 1.30	943.25	–	0.2586	0.05
HA2	int.	18.40	3.24 ± 0.04	2.56 ± 0.14	1.00 ± 0.20	0.69 ± 0.03	0.63 ± 1.37	182.59 ± 68.77	0.66 ± 0.09	5.80 ± 1.38	1.56 ± 0.81	919.85	4.16	0.1017	0.05
HA3	int.	18.40	2.79 ± 0.05	2.37 ± 0.14	1.11 ± 0.24	0.68 ± 0.03	0.70 ± 1.36	171.36 ± 59.27	0.66 ± 0.08	6.19 ± 1.40	1.55 ± 0.76	942.35	3.98	0.1027	0.05
HA6	int.	11.35	1.61 ± 0.05	2.52 ± 0.13	0.97 ± 0.19	0.69 ± 0.03	0.74 ± 1.85	182.32 ± 60.37	0.67 ± 0.08	7.12 ± 1.82	1.60 ± 0.89	922.05	3.18	0.0992	0.05
IA0	int.	14.45	3.61 ± 0.03	4.21 ± 0.17	0.65 ± 0.06	0.74 ± 0.01	1.30 ± 1.48	342.25 ± 109.72	0.50 ± 0.10	4.79 ± 1.90	1.86 ± 1.33	455.54	0.81	0.0856	0.05
IA1	int.	11.35	1.82 ± 0.03	2.96 ± 0.59	0.58 ± 0.09	0.77 ± 0.04	0.36 ± 0.46	141.89 ± 210.99	0.55 ± 0.45	1.57 ± 0.16	1.66 ± 0.48	606.39	35.15	0.0958	0.04
IA2	int.	18.40	3.57 ± 0.02	3.71 ± 0.18	0.63 ± 0.07	0.75 ± 0.02	0.96 ± 0.92	316.93 ± 120.33	0.50 ± 0.11	4.39 ± 1.02	2.12 ± 1.08	519.98	1.73	0.0752	0.04
IA3	int.	14.45	2.40 ± 0.03	2.83 ± 0.43	0.54 ± 0.10	0.78 ± 0.03	0.50 ± 0.23	191.14 ± 175.30	0.51 ± 0.29	1.82 ± 0.32	1.60 ± 0.59	617.41	15.54	0.0995	0.05
IA6	int.	11.35	1.61 ± 0.03	4.26 ± 0.33	0.48 ± 0.05	0.77 ± 0.02	0.42 ± 0.16	248.66 ± 233.70	0.54 ± 0.28	2.12 ± 0.39	1.78 ± 0.84	470.16	10.69	0.0894	0.05
CA0	cont.	11.35	11.93 ± 0.05	1.81 ± 0.21	1.81 ± 0.48	0.61 ± 0.04	1.99 ± 0.50	311.20 ± 57.84	0.38 ± 0.06	5.13 ± 0.82	2.36 ± 0.67	2004.98	0.56	0.0674	0.03
CA1	cont.	14.45	6.10 ± 0.03	1.36 ± 0.10	1.03 ± 0.23	0.72 ± 0.03	1.24 ± 0.72	371.95 ± 61.15	0.44 ± 0.05	4.81 ± 0.82	1.76 ± 0.62	1506.00	0.93	0.0906	0.03
CA2	cont.	23.44	7.20 ± 0.02	1.86 ± 0.08	0.95 ± 0.13	0.70 ± 0.02	1.48 ± 1.23	350.49 ± 50.18	0.47 ± 0.04	7.23 ± 1.11	2.01 ± 0.73	1371.93	0.64	0.0790	0.03
CA3	cont.	18.40	4.78 ± 0.02	2.06 ± 0.06	0.57 ± 0.07	0.75 ± 0.02	2.35 ± 3.23	407.56 ± 44.21	0.47 ± 0.04	7.60 ± 4.45	1.09 ± 1.01	1222.92	0.17	0.1465	0.03
CA3W	cont.	23.44	5.09 ± 0.02	2.11 ± 0.06	0.55 ± 0.07	0.75 ± 0.02	2.24 ± 2.64	389.60 ± 43.30	0.46 ± 0.04	7.10 ± 3.66	1.12 ± 0.92	1235.72	0.21	0.1419	0.03
CA6	cont.	14.45	3.10 ± 0.03	1.51 ± 0.04	0.36 ± 0.06	0.82 ± 0.02	1.67 ± 2.84	432.84 ± 28.13	0.51 ± 0.02	0.70 ± 7.40	0.10 ± 1.07	1453.43	0.30	–	0.04

Table 3.4: Finite open Warburg model fitting, cathode.

ID	loading	HF_cut	Rs	R1	Q1	n1	R2	Q2	n2	Rw	tauD	f_c1	f_c2	f_W	RMS
unit		kHz	Ω	Ω	m $\frac{s^n}{\Omega}$	-	Ω	m $\frac{s^n}{\Omega}$	-	Ω	s	Hz	Hz	Hz	-
UC0	-	29.85	10.15 ± 0.03	2.61 ± 0.11	0.67 ± 0.09	0.69 ± 0.02	1.84 ± 0.78	256.71 ± 33.52	0.45 ± 0.04	6.98 ± 0.96	2.08 ± 0.55	1603.29	0.84	0.0766	0.03
UC1	-	23.44	14.28 ± 0.02	7.33 ± 0.07	0.25 ± 0.01	0.78 ± 0.01	4.75 ± 30.10	289.15 ± 10.34	0.52 ± 0.01	0.15 ± 88.88	0.03 ± 20.93	520.85	0.09	–	0.05
UC2	-	29.85	14.25 ± 0.02	8.35 ± 0.07	0.22 ± 0.01	0.79 ± 0.01	5.79 ± 17.97	232.87 ± 8.45	0.50 ± 0.02	0.22 ± 52.70	0.05 ± 12.67	470.29	0.09	–	0.06
UC3	-	8.91	14.30 ± 0.04	3.21 ± 0.22	0.30 ± 0.04	0.77 ± 0.02	2.79 ± 0.14	139.96 ± 19.30	0.36 ± 0.04	5.61 ± 0.54	1.91 ± 0.28	1282.97	2.17	0.0835	0.04
UC3W	-	11.35	12.67 ± 0.03	4.26 ± 0.12	0.42 ± 0.03	0.73 ± 0.01	1.44 ± 0.58	233.37 ± 36.07	0.49 ± 0.05	5.90 ± 0.76	2.07 ± 0.51	960.45	1.45	0.0771	0.03
UC6	-	11.35	16.88 ± 0.03	4.40 ± 0.12	0.36 ± 0.03	0.74 ± 0.01	1.10 ± 0.42	257.68 ± 47.54	0.50 ± 0.06	3.11 ± 0.81	1.23 ± 0.48	955.11	1.97	0.1298	0.04
HC0	int.	14.45	13.06 ± 0.05	5.83 ± 0.12	0.59 ± 0.05	0.64 ± 0.01	1.63 ± 1.01	105.52 ± 12.66	0.67 ± 0.03	6.56 ± 1.76	0.80 ± 0.30	1100.84	2.16	0.1978	0.05
HC1	int.	14.45	15.13 ± 0.21	4.35 ± 0.32	1.84 ± 0.40	0.48 ± 0.03	1.42 ± 0.62	78.87 ± 10.79	0.70 ± 0.03	5.44 ± 1.25	0.68 ± 0.21	3705.67	3.61	0.2352	0.05
HC2	int.	18.40	11.74 ± 0.09	4.25 ± 0.17	1.20 ± 0.19	0.57 ± 0.02	1.07 ± 1.68	90.81 ± 18.23	0.74 ± 0.04	8.45 ± 2.28	0.88 ± 0.40	1766.16	3.66	0.1815	0.05
HC3	int.	18.40	14.55 ± 0.10	4.24 ± 0.20	1.21 ± 0.21	0.56 ± 0.02	1.21 ± 0.72	89.20 ± 14.33	0.69 ± 0.04	6.87 ± 1.14	1.00 ± 0.27	1924.22	4.05	0.1595	0.05
HC6	int.	11.35	14.95 ± 0.29	4.41 ± 0.43	2.85 ± 0.68	0.43 ± 0.04	0.98 ± 0.89	84.98 ± 18.22	0.75 ± 0.04	5.99 ± 1.48	0.78 ± 0.30	3897.33	4.39	0.2042	0.05
IC0	int.	14.45	15.21 ± 0.03	5.75 ± 0.14	0.41 ± 0.03	0.69 ± 0.01	1.88 ± 0.58	176.02 ± 22.88	0.50 ± 0.04	6.29 ± 0.93	1.64 ± 0.40	981.02	1.47	0.0971	0.04
IC1	int.	11.35	20.54 ± 0.04	11.89 ± 0.10	0.28 ± 0.01	0.72 ± 0.01	7.06 ± 16.04	164.35 ± 5.19	0.52 ± 0.01	0.36 ± 46.76	0.06 ± 7.74	436.82	0.12	–	0.06
IC2	int.	18.40	13.62 ± 0.03	9.43 ± 0.08	0.33 ± 0.01	0.72 ± 0.00	3.18 ± 1.94	213.08 ± 19.54	0.62 ± 0.02	4.64 ± 4.21	0.97 ± 0.99	488.36	0.30	0.1633	0.05
IC3	int.	14.45	14.00 ± 0.03	8.07 ± 0.08	0.35 ± 0.01	0.72 ± 0.01	2.42 ± 1.30	223.12 ± 23.14	0.57 ± 0.03	6.10 ± 2.37	1.41 ± 0.71	571.99	0.47	0.1131	0.04
IC6	int.	11.35	18.68 ± 0.04	7.49 ± 0.12	0.33 ± 0.02	0.71 ± 0.01	2.84 ± 0.82	136.22 ± 12.25	0.55 ± 0.03	5.17 ± 1.77	0.95 ± 0.39	767.02	0.89	0.1679	0.05
CC0	cont.	11.35	15.66 ± 0.03	4.03 ± 0.10	0.39 ± 0.04	0.74 ± 0.01	1.65 ± 0.46	217.66 ± 26.68	0.50 ± 0.03	6.19 ± 0.71	2.52 ± 0.49	970.65	1.22	0.0632	0.04
CC1	cont.	14.45	14.51 ± 0.03	4.17 ± 0.10	0.36 ± 0.03	0.75 ± 0.01	2.00 ± 0.45	200.46 ± 20.55	0.48 ± 0.03	4.80 ± 0.90	1.57 ± 0.40	927.88	1.05	0.1017	0.04
CC2	cont.	23.44	11.79 ± 0.02	5.04 ± 0.09	0.31 ± 0.02	0.76 ± 0.01	2.78 ± 1.93	199.34 ± 21.26	0.53 ± 0.03	10.16 ± 2.71	1.63 ± 0.69	750.31	0.49	0.0977	0.05
CC3	cont.	18.40	15.95 ± 0.03	6.45 ± 0.10	0.27 ± 0.02	0.77 ± 0.01	2.73 ± 1.57	180.10 ± 19.37	0.54 ± 0.03	9.02 ± 2.44	1.56 ± 0.62	601.79	0.59	0.1021	0.05
CC3W	cont.	23.44	15.37 ± 0.03	6.59 ± 0.11	0.37 ± 0.03	0.73 ± 0.01	2.89 ± 2.99	221.65 ± 30.82	0.59 ± 0.04	8.08 ± 5.22	1.45 ± 1.23	601.72	0.34	0.1096	0.06
CC6	cont.	14.45	17.26 ± 0.10	3.20 ± 0.24	2.18 ± 0.49	0.53 ± 0.03	2.10 ± 1.06	176.63 ± 26.87	0.53 ± 0.04	8.18 ± 1.48	1.73 ± 0.55	2001.71	1.02	0.0918	0.04

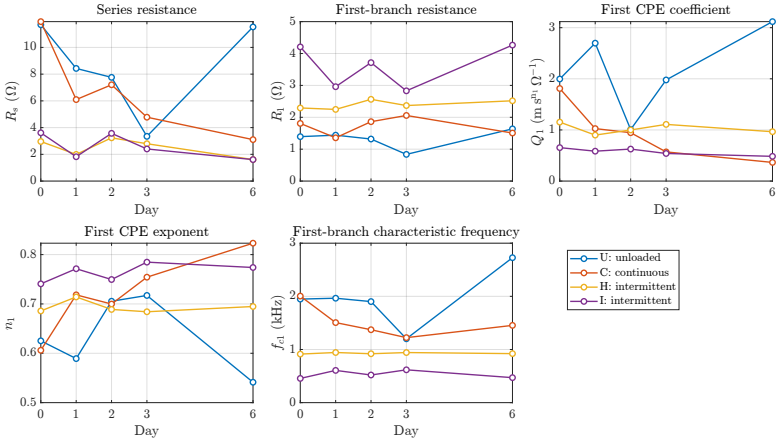


Figure 3.8: Temporal evolution of selected fitted anode parameters for the unloaded cell (U), continuously loaded cell (C), and intermittently loaded cells (H and I). The panels show the series resistance R_s, first-branch resistance R_1, first-CPE coefficient Q_1, first-CPE exponent n_1, and first-branch characteristic frequency f_{c1}.

3.3.2 Effect of electrical loading

In this subsection, we compare the response of both anodes and cathodes under different electrical loading conditions. First, we compare the most robust fitted parameters among the cells and examine their evolution over time, before interpreting the main loading-related trends.

Anode response to loading: Figure 3.8 summarizes the temporal evolution of selected fitted anode parameters describing the ohmic and first-interfacial response, namely R_s, R_1, Q_1, n_1, and f_{c1}, for the unloaded, continuously loaded, and intermittently loaded cells. The Day 3W measurements are excluded because they correspond to the immediate post-watering condition considered separately in Sec. 3.3.3.

Among the tested cells, IA exhibits the highest R_1 values, ranging from 2.8 to 4.3 Ω, followed by HA at 2.3–2.6 Ω, CA at 1.4–2.1 Ω, and UA at 0.8–1.6 Ω. The R_1 values of HA remain relatively stable throughout the test, while IA varies without a clear monotonic trend and remains higher than the other cells. In contrast, CA shows a slight increase from 1.8 to 2.1 Ω, whereas UA decreases from 1.4

to 0.9 Ω. After Day 3, the two cells diverged: CA decreased to 1.5 Ω by Day 6, whereas UA increased to 1.6 Ω. At the same time, IA and HA exhibit relatively stable n_1 values, ranging from 0.74 to 0.78 and from 0.68 to 0.71, respectively. In contrast, CA and UA show an increasing trend during the first three days of testing, with n_1 increasing from 0.61 to 0.75 and from 0.63 to 0.72, respectively. During the last three days of testing, the trends of CA and UA diverge: n_1 increases further to 0.82 for CA, whereas it decreases to 0.54 for UA.

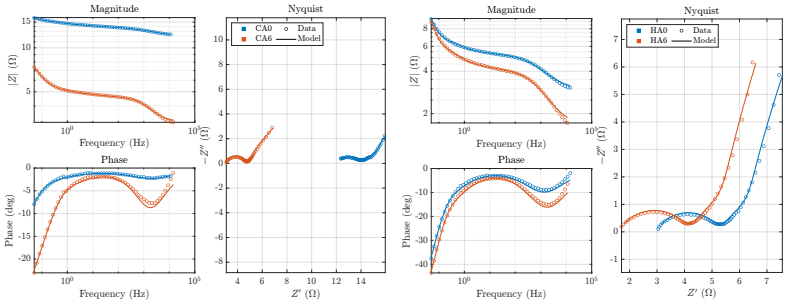
The Q_1 values remain relatively stable for IA and HA, ranging from 0.5 to 0.7 $\mathrm{m\,s}^n\,\Omega^{-1}$ and from 0.9 to 1.2 $\mathrm{m\,s}^n\,\Omega^{-1}$, respectively. In contrast, UA varies non-monotonically, remaining close to its initial value on Day 3 before increasing to 3.1 $\mathrm{m\,s^n\Omega^{-1}}$ on Day 6. CA shows the opposite behaviour, with a continuous decrease throughout the test from 1.81 to 0.36 $\mathrm{m\,s}^n\,\Omega^{-1}$. Furthermore, the group-averaged R_s values followed the order

$$\begin{aligned} &\overline{R}_{s,\mathrm{HA}} = 2.52\ \Omega\ < \overline{R}_{s,\mathrm{IA}} = 2.60\ \Omega \\ &< \overline{R}_{s,\mathrm{CA}} = 6.62\ \Omega\ < \overline{R}_{s,\mathrm{UA}} = 8.55\ \Omega \end{aligned} \tag{3.5}$$

Similarly to the other parameters, the R_s values of IA and HA remain relatively stable, ranging from 1.6 to 3.6 Ω and from 1.6 to 3.2 Ω, respectively. In contrast, UA exhibits a marked decrease during the first three days of testing, from 11.7 to 3.35 Ω, followed by an increase during the last three days, reaching 11.5 Ω. CA shows an overall decrease from 11.9 Ω on Day 0 to 3.1 Ω on Day 6. Ultimately, these parameter trends are reflected in the characteristic frequency of the first interfacial process. The group-averaged f_{c1} values increase in the order IA < HA < CA < UA:

$$\begin{aligned} &\overline{f}_{c1,\mathrm{IA}} = 534\ \mathrm{Hz} < \overline{f}_{c1,\mathrm{HA}} = 928\ \mathrm{Hz} \\ &< \overline{f}_{c1,\mathrm{CA}} = 1512\ \mathrm{Hz} < \overline{f}_{c1,\mathrm{UA}} = 1949\ \mathrm{Hz}. \end{aligned} \tag{3.6}$$

Thus, IA exhibits the slowest apparent first interfacial process, whereas UA exhibits the fastest one. The f_{c1} values remain relatively stable for IA and HA, around 534 Hz and 928 Hz, respectively, while CA decreases from 2005 Hz on Day 0 to 1223 Hz on Day 3 and then partially recovers to 1453 Hz on Day 6. UA displays a change in trend: during the first three days, f_{c1} decreases from 1948 to 1201 Hz, whereas during the last three days it increases from 1201 to 2727 Hz.



(a) Continuous loading, Cell C. (b) Intermittent loading, Cell H.

Figure 3.9: Representative anode EIS fits showing the temporal evolution between Day 0 and Day 6 under continuous and intermittent loading. The continuously loaded anode shows a stronger temporal shift, whereas the intermittently loaded anode remains comparatively closer to its initial response.

The representative anode fits in Fig. 3.9 illustrate the temporal evolution behind these parameter trends. The continuously loaded anode shows a strong displacement between Day 0 and Day 6, consistent with a progressive modification of the anodic interfacial response during sustained polarization. The intermittently loaded Cell H also evolves, but its displacement is smaller and remains within the same general spectral shape. This supports the interpretation that the clearest loading-dependent changes occur at the anode, while the intermittent profile does not produce the same progressive shift observed under continuous loading.

Overall, both intermittently operated cells preserve a comparatively stable biofilm/electrode interface throughout the test. IA and HA exhibit the lowest f_{c1} and R_s, as well as the lowest group-averaged Q_1 values, together with the highest n_1 and R_1 values. This indicates a slower apparent biofilm/electrode interfacial process, a lower ohmic contribution from the electrolyte and electrode/-electrolyte contact, and a more resistive biofilm/electrode interface. At the same time, the reduced Q_1 values and higher n_1 values suggest a lower apparent interfacial storage contribution and a more homogeneous, closer-to-ideal double-layer capacitive response.

Clear differences emerge when comparing the continuously loaded anode and the unloaded anode. During the first three days, CA and UA show similar trends for most parameters, except for the weakly

opposite behaviour observed in R_1. Their Q_1 responses differ: CA decreases progressively, whereas UA shows a nonmonotonic evolution during the first three days and returns close to its initial value on Day 3. In both cases, the increase in n_1, together with the decrease in R_s and f_{c1}, suggests that the interfacial response shifted toward longer timescales while becoming more homogeneous. For CA, the decrease in Q_1 may indicate a reduction in the apparent fast interfacial storage contribution, possibly associated with changes in active biofilm area, hydration state, or charge-storage distribution.

The behaviour of UA changes during the last three days, when no EIS measurements were performed. During the first three days, CA and UA show broadly similar trends for R_s, f_{c1}, and n_1. After Day 3, however, their behaviour diverges: CA continues to decrease in R_s and Q_1 and increase in n_1, although f_{c1} partially recovers by Day 6, whereas UA reverses its trends. Specifically, UA shifts toward higher ohmic resistance, a larger apparent interfacial storage contribution, a faster interfacial process, and a less homogeneous interface. This change in trend connected to the suspension of the EIS measurements suggests that, even in the absence of external loading, the repeated EIS measurements performed during the first three days may have weakly perturbed or conditioned the unloaded cell in a direction similar to continuous electrical operation. Once the EIS measurements were halted during the last three days, UA therefore evolved toward a different state, closer to its undisturbed unloaded behaviour. Interestingly, the R_1 trends are opposite for UA and CA during both the first and last three days. However, since this variation is less pronounced than the trends observed for the other parameters, it is not considered sufficient to undermine the overall interpretation.

Thus, the closest representation of the true unloaded state appears to correspond to UA0 and UA6. These conditions are characterized by a faster interfacial process, a more inhomogeneous interface, higher ohmic resistance, and a larger apparent interfacial storage contribution. Conversely, continuous loading appears to slow down the interfacial process, make the interface more homogeneous, and reduce both the ohmic resistance and the apparent interfacial storage contribution. Compared with CA and UA, HA and IA display a more stable behaviour over time. They are characterized by the slowest first-branch interfacial process, the highest charge-transfer resistance, and the lowest ohmic and interfacial stor-

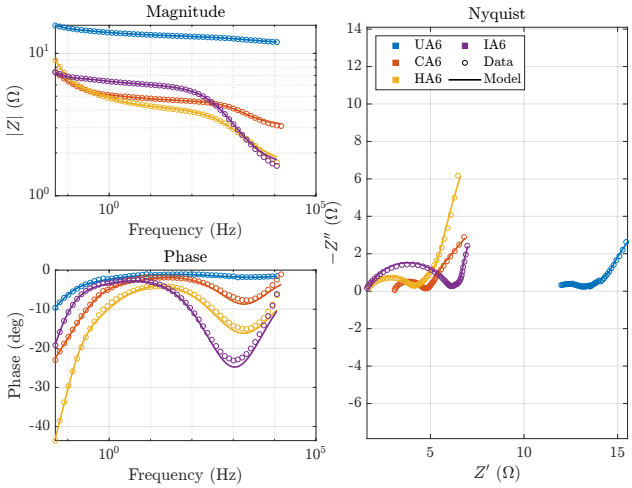


Figure 3.10: Representative late-state anode EIS fits for unloaded, continuously loaded, and intermittently loaded cells. The comparison illustrates the effect of loading history on the anodic impedance response at Day 6.

age contributions. The distinct anode responses observed under different loading histories are consistent with three-electrode studies showing that polarization history affects electroactive-biofilm behaviour. These results provide promising evidence that comparable impedance variations also arise under two-electrode harvesting; future work should validate their statistical significance and long-term evolution using larger PMFC populations and extended monitoring.

The late-state comparison in Fig. 3.10 further shows that the final anodic spectra remain distinguishable according to loading history. The unloaded anode retains a clearly different high-impedance response, while the continuously and intermittently loaded anodes occupy lower-impedance regions with distinct spectral shapes. This supports the conclusion that electrical loading leaves a measurable signature in the anodic impedance response. Cathode loading-history fits are not shown here because the cathode did not exhibit an equally ordered loading-dependent trend and was more strongly affected by hydration and local contact conditions.

Although some second-branch model parameters appear to show temporal trends, such as f_W in CA and R_2 in UA, these trends

are not considered conclusive because of the high standard errors associated with the limited identifiability of the diffusion-related branch.

Cathode response to loading: Across the cathode measurements, R_s ranged from 10.15 to 16.88 Ω for UC, from 11.74 to 15.13 Ω for HC, from 13.62 to 20.54 Ω for IC, and from 11.79 to 17.26 Ω for CC. Similarly, R_1 ranged from 2.61 to 8.35 Ω, 4.24 to 5.83 Ω, 5.75 to 11.89 Ω, and 3.20 to 6.59 Ω for UC, HC, IC, and CC, respectively. These cathode R_s and R_1 values were generally higher than the corresponding anode values, indicating a larger ohmic/contact resistance and a larger oxygen-reduction/interfacial limitation at the cathode. The cathode Q_1 values ranged from 0.22 to 0.67 $\mathrm{m\,s^{n}\,\Omega^{-1}}$ for UC, 0.59 to 2.85 $\mathrm{m\,s^{n}\,\Omega^{-1}}$ for HC, 0.28 to 0.41 $\mathrm{m\,s^{n}\,\Omega^{-1}}$ for IC, and 0.27 to 2.18 $\mathrm{m\,s^{n}\,\Omega^{-1}}$ for CC, suggesting a generally lower or comparable apparent double-layer or pseudocapacitive contribution in the main interfacial branch. The remaining fitted parameters were of comparable magnitude to those obtained for the anode. Unlike the anode, the cathode did not show a clear monotonic loading-induced evolution across time or between loading conditions. Therefore, the cathode response was interpreted as being more strongly influenced by hydration than by the applied loading regime alone.

3.3.3 Effect of hydration

Although the original aim of the study was to investigate loading effects, the water content varied unpredictably during operation. In particular, the intermittently loaded cells appeared to dry out more rapidly than the continuously loaded and unloaded cells. Therefore, humidity-related effects had to be considered explicitly in the interpretation of the EIS measurements. In addition, the immediate response to water replenishment was analysed in order to distinguish effects that appeared directly after wetting from those that developed over longer time scales, as well as to identify the frequency regions most affected by hydration.

Hydration effects were therefore interpreted separately from loading effects by considering two characteristic time scales: the immediate Day 3-Day 3W response for the continuously loaded and unloaded cells C and U, and the Day 1-Day 2 response following the visibly dry condition observed in the intermittently loaded cells H

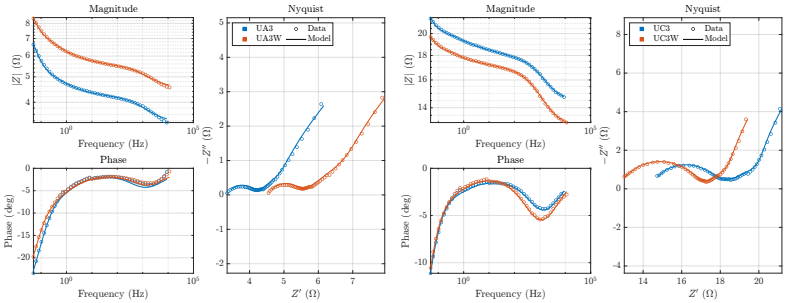
and I.

Cells I and H were visibly dry on Day 1, whereas cells C and U never exhibited an evident dry state. Therefore, for cells I and H, the Day 1 measurements were taken as representative of the clearest dry condition, while for cells C and U the analysis focused on the measurements furthest from the watering events, namely Day 3 and Day 6. We also considered Day 6 measurements for cells I and H, for completeness.

For the anodes, the lowest R_s values within each cell were often observed either in measurements performed further from watering events, namely UA3 ($R_s = 3.6\ \Omega$), CA3 ($R_s = 4.8\ \Omega$), IA6 ($R_s = 1.6\ \Omega$), and HA6 ($R_s = 1.6\ \Omega$), or under the visibly driest conditions, namely IA1 ($R_s = 2.0\ \Omega$) and HA1 ($R_s = 1.8\ \Omega$). This behaviour is not universal, as shown by the high R_s value of UA6 ($R_s = 11.5\ \Omega$). In the clearest dry states, IA1 and HA1, the Warburg-related resistance and diffusion time constant were also minimized, with IA1 showing $R_W = 1.6\ \Omega$ and $\tau_D = 1.7$ s, and HA1 showing $R_W = 4.0\ \Omega$ and $\tau_D = 0.6$ s. A similar behaviour was observed for UA3, where $R_W = 4.7\ \Omega$ and $\tau_D = 1.0$ s. This trend was not observed for CA3, possibly because this cell retained a higher hydration level than UA3.

In contrast, the cathode-side measurements IC1 and HC1 showed the highest R_s values, equal to 20.54 Ω and 15.13 Ω, respectively. In the case of IC1, R_1 also reached its maximum value, 11.9 Ω. This pronounced increase in R_1 may be related to cell I representing the driest observed condition, which could have enhanced the cathode interfacial limitation. Overall, this suggests that the hydration state affected the anode and cathode responses differently, with drying being associated with reduced anode R_s and diffusion-related contributions, but increased cathode ohmic/contact and interfacial resistances.

Immediate watering response: Day 3 to Day 3W, cells C and U: The immediate effect of watering is illustrated in Fig. 3.11. The response is electrode-dependent and, unexpectedly, watering produced only limited changes in the fitted parameters. At the anode, wetting did not simply reduce all resistive contributions. In UA, R_s increased slightly from 3.35 to 4.54 Ω, while R_1 changed only marginally from 0.84 to 0.91 Ω. The main variation was observed in the low-frequency branch, where R_W increased from 4.69 to 8.83 Ω, and τ_D increased from 1.00 to 2.25 s. This suggests that the short-



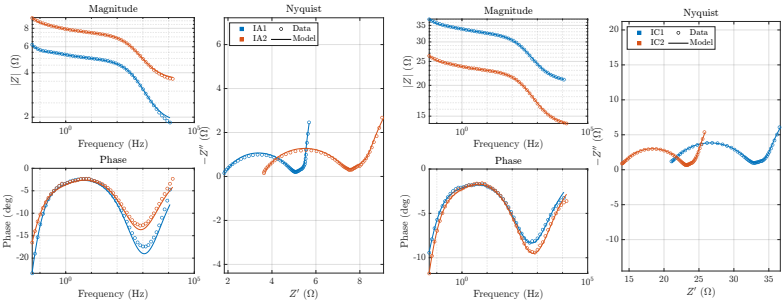
(a) Unloaded anode. (b) Unloaded cathode.

Figure 3.11: Representative EIS fits showing the immediate effect of watering between Day 3 and Day 3W for the unloaded cell. The anode and cathode respond differently, highlighting the electrode-specific impact of hydration on the fitted impedance spectra.

term wetting perturbation mainly affected the apparent bounded-transport response of UA, rather than the high-frequency ohmic or interfacial terms. In CA, immediate wetting had little effect, as all fitted parameters changed only slightly from CA3 to CA3W, likely because this cell had not reached a strongly perturbed dry state before wetting.

On the cathode side, wetting produced a clearer change in the ohmic/contact contribution. In UC, R_s decreased from 14.30 to 12.67 Ω, suggesting improved ionic or contact conductivity. However, R_1 increased from 3.21 to 4.26 Ω, indicating that wetting did not immediately improve the oxygen-reduction/interfacial process. CC behaved similarly to UC, with only minor changes in both R_s and R_1. Overall, immediate watering appears to affect mainly the low-frequency transport-related response at the anode, while slightly reducing the cathodic ohmic/contact resistance. The simultaneous marginal increase in cathode R_1 should not be interpreted as a persistent hydration trend, since the Day 1–Day 2 post-watering comparison of the intermittently loaded cathodes shows either a decrease, as in IC, or only a small change, as in HC. It is therefore more likely a transient interfacial effect associated with pore filling, ion redistribution, changes in contact geometry, oxygen availability, or local redox conditions.

Post-watering response after the dry Day 1 condition, cells I and H: the Day 1–Day 2 transition was particularly relevant



(a) Cell I anode. (b) Cell I cathode.

Figure 3.12: Representative EIS fits showing the delayed post-watering response of Cell I between Day 1 and Day 2. The anode response shows changes in both the high-frequency/contact region and the low-frequency bounded-transport region. The cathode response shows a shift of the high-frequency intercept and first interfacial arc toward lower real impedance after watering.

for the intermittently operated cells, as water had to be replenished after the visibly dry Day 1 measurement. A delayed post-watering response is shown in Fig. 3.12 for Cell I. In this case, both electrodes show a marked displacement between Day 1 and Day 2, but with different spectral signatures. At the anode, water replenishment mainly restored the low-frequency bounded-transport contribution in both intermittent cells, while also affecting the high-frequency/contact region, especially in the visibly drier Cell I. In HA, one day after watering, R_W increased from 4.0 to 5.8 Ω, while τ_D increased from 0.62 to 1.56 s. These values approached those measured before drying in HA0 ($R_W = 6.1$ Ω, $\tau_D = 1.4$ s), suggesting a partial restoration of the bounded-transport response. At the same time, HA also showed an increase in R_s, from 1.99 to 3.24 Ω, and a moderate increase in R_1, from 2.25 to 2.56 Ω, although f_{c1} remained almost unchanged, from 943 to 920 Hz. Therefore, the high-frequency change in HA appears mainly as an ohmic/contact shift rather than as a strong modification of the first interfacial process.

In IA, the post-watering change was stronger. Watering increased R_W from 1.6 to 4.4 Ω and τ_D from 1.7 to 2.1 s, bringing the IA2 values close to those of IA0 ($R_W = 4.8$ Ω, $\tau_D = 1.9$ s). In addition, IA showed a larger high-frequency displacement, with R_s

increasing from 1.82 to 3.57 Ω, R_1 increasing from 2.96 to 3.71 Ω, and f_{c1} decreasing from 606 to 520 Hz. This indicates that the dry Day 1 condition in Cell I affected not only the slower transport-related response, but also the near-electrode/contact and first interfacial contribution. Overall, the visibly dry Day 1 condition was associated with reduced R_W and τ_D, whereas after watering these parameters recovered toward their Day 0 values over the following day. However, the simultaneous changes in R_s and, more clearly for IA, in R_1 show that hydration effects cannot be confined exclusively to the low-frequency branch. The effect cannot be attributed exclusively to watering because one additional day of operation also occurred between the two measurements.

At the cathode, the post-watering response is stronger in the high-frequency and first interfacial region, with a marked reduction in the real part of the impedance. In HC, R_s decreased from 15.1 to 11.7 Ω, while R_1 changed only slightly. In IC, which was visibly drier, R_s decreased from 20.5 to 13.6 Ω, while R_1 decreased from 11.9 to 9.4 Ω. The R_s values after watering approached the corresponding Day 0 values of IC0 ($R_s = 15.1$ Ω) and HC0 ($R_s = 13.1$ Ω), indicating that water replenishment partially restored the cathodic ohmic/contact contribution. In contrast, R_1 recovered only partially, especially in IC, whose Day 0 value was $R_1 = 5.7$ Ω. This suggests that the cathodic response is strongly influenced by hydration, local contact conditions, and oxygen access, and that the recovery of the oxygen-reduction/interfacial process can be slower or less complete than the recovery of the ohmic contribution.

The representative fits therefore support the decision to treat hydration as a major confounding factor when interpreting loading-dependent EIS trends. Poorer identifiability of the R_2 and Q_2 parameters was observed in the measurements furthest from watering events or corresponding to the driest states, namely HA1, IA1, IA3, IA6, and especially IC1. This suggests that, under dry or poorly hydrated conditions, the second R–CPE branch may have partially absorbed part of the diffusion-like response, making the low-frequency model structure less well resolved. At the same time, the HA1–HA2 and IA1–IA2 comparisons show that drying and subsequent water replenishment can also perturb the high-frequency/contact region of the anode, especially through R_s. This effect is stronger in IA, where R_1 and f_{c1} also change more clearly, indicating that severe drying can affect the first interfacial branch as well. Hydration ef-

fects therefore cannot be confined exclusively to the low-frequency branch. The effect was particularly evident for Cell I: after the dry condition reached on Day 1, satisfactory identifiability was recovered only in IA2 following water replenishment, whereas IA3 and IA6 again showed identifiability problems.

3.3.4 Impedance-analysis summary and outlook

The same impedance model was successfully fitted to both anode and cathode spectra under the different operating conditions. Overall, the results suggest that electrical loading primarily affected the high-frequency anodic biofilm/electrode response. Hydration mainly influenced the low-frequency anodic bounded-transport response and the high-frequency cathodic ohmic/interfacial response, but the HA1–HA2 and IA1–IA2 comparisons also show that severe drying and subsequent water replenishment can perturb the high-frequency/contact region of the anode, especially through R_s and, in IA, also through R_1 and f_{c1}. The first R–CPE branch provided the most reliable comparison between electrodes and operating modes, while the second R–CPE branch and finite Warburg element were useful mainly as qualitative indicators of drying and transport-related changes because of stronger parameter coupling and limited low-frequency identifiability. These results show that loading-history-dependent electrode responses can be tracked by EIS at polarization points derived from the harvesting state, providing an electrode-resolved counterpart to the full-cell two-electrode operation. The observed trends are consistent with the loading-history effects reported under controlled three-electrode polarization, while longer experiments and larger PMFC populations are needed to establish their robustness and persistence. This harvesting-referenced characterization is particularly relevant for PMFC-powered embedded nodes, where prolonged unloading to recover open-circuit conditions would interrupt energy extraction and may not represent the electrochemical state experienced during normal harvesting operation.

The main effect of electrical loading was observed at the anode, especially in the first R–CPE branch. IA and HA showed the most stable anode parameters over time, suggesting that intermittent loading preserved a comparatively stable biofilm/electrode interface. In particular, the intermittent anodes showed the lowest

first-branch characteristic frequency and the lowest ohmic contribution, together with the highest charge-transfer resistance. Their Q_1 values were also the lowest on average across the corresponding groups. These effects could be caused by the intermittent loading treatment and thus offer a promising research direction for future work.

On the other hand, CA showed clear trends over time, with decreasing R_s and Q_1, f_{c1} decreasing through Day 3 and partially recovering by Day 6, as well as increasing n_1, indicating a progressive conditioning of the anodic interface (less ohmic resistance and interfacial storage contribution, a slower interfacial process and a more homogeneous interface) during continuous operation. UA followed a similar trend during the first part of the test, suggesting that repeated EIS measurements may also condition the anode weakly. This is supported by the fact that after the interval without EIS measurements, UA returned to a situation which we deem representative of truly unloaded state, with more ohmic resistance, interfacial storage contribution, a faster interfacial process and a less homogeneous interface. Cathode R_s and R_1 were generally higher than the corresponding anode values, confirming that the cathode contributed a larger ohmic/contact and ORR/interfacial limitation. However, cathode trends were not clearly ordered by loading condition. The strongest cathode perturbations occurred in the driest states, especially IC, indicating that hydration, oxygen access, and local contact conditions dominated the cathode response more than loading alone.

Hydration affected the anode and cathode differently. The hydration analysis was used to identify which parts of the spectrum were most sensitive to drying and whether water replenishment restored the fitted parameters after visible drying. The Day 1 measurements of the intermittently loaded cells were treated as dry-state references, especially for Cell I, while the Day 3/Day 3W comparisons were used to assess the immediate effect of watering in the unloaded and continuously loaded cells.

For the anodes, the lowest R_s values within each cell were often observed either in measurements performed further from watering events, namely UA3, CA3, IA6, and HA6, or under the visibly driest conditions, namely IA1 and HA1. This behaviour is not universal, as shown by the high R_s value of UA6. Therefore, the decrease in fitted anodic R_s under dry or far-from-watering conditions may

reflect changes in local ionic concentration, electrode–soil contact geometry, or redistribution of the fitted contributions between R_s and the first interfacial branch.

In the clearest dry states, IA1 and HA1, the Warburg-related resistance and diffusion time constant were minimized. A similar behaviour was observed for UA3, but not for CA3, possibly because CA retained a higher hydration level. Drying also reduced the identifiability of the second R–CPE branch, especially in IA1, IC1, and HA1, suggesting that the slow interfacial and transport-related responses became less clearly separable in the fitted model.

In contrast, cathode drying manifested mainly as an increase in the ohmic/contact and first interfacial terms. The cathode-side measurements IC1 and HC1 showed the highest R_s values in their groups, and IC1 also showed the highest R_1. Delayed post-watering measurements on IA and HA showed that water replenishment restored the bounded-transport contribution toward the pre-drying values, while also perturbing the high-frequency/contact region of the anode, especially in IA. Immediate post-watering measurements on UA and CA showed a clear restoration only where a drying perturbation was present, namely in UA.

For the cathodes, hydration mainly restored the high-frequency resistance terms. In the immediate post-watering measurement, UC and CC showed reduced R_s and slightly increased R_1, suggesting a transient interfacial effect rather than a persistent hydration trend. In the delayed measurement, watering reduced R_s in IC and HC, indicating improved ionic/contact conductivity, and in IC it also reduced R_1. However, while the cathodic ohmic resistance was largely restored after the strong dry-state perturbation, the oxygen-reduction/interfacial resistance R_1 recovered only partially, especially in IC.

Future work should extend the impedance measurements to lower frequencies to better resolve the second impedance branch and the finite-diffusion contribution. This extension should also reassess the lower-frequency bound, identifying whether it can be shifted while still satisfying the time and energy constraints of battery-free EIS operation. The present results suggest that, especially under drier conditions, the finite Warburg element and the second R–CPE branch may partially absorb each other under the given model, making the fitted low-frequency parameters less identifiable for both anodes and cathodes. Extending the low-frequency range would help

distinguish these contributions more clearly. However, this may be challenging in a battery-free measurement context, because lower-frequency EIS points require longer acquisition times and higher active energy consumption. Finally, the trends identified in this preliminary study indicate the need for a larger experimental setup including biological replicates, temperature and humidity control, and direct tracking of the cell water content. Therefore, this work should be considered a preliminary investigation, whose results provide a basis and motivation for a broader study aimed at validating the observed loading- and hydration-related trends under more controlled and statistically robust conditions. Once these impedance trends are validated, they could be used as feedback variables to adapt the harvesting voltage, period, and duty cycle to the evolving PMFC state.

3.4 Conclusions

This chapter addressed the problem of using PMFCs as practical energy sources for battery-free embedded systems. PMFCs offer a sustainable route to long-term low-power generation, but their low voltage, low power density, slow dynamics, and dependence on biological and environmental conditions make conventional harvesting and monitoring strategies difficult to apply directly. For this reason, the chapter focused on two coupled aspects of PMFC energy management: controlled intermittent harvesting and non-destructive impedance-based state observation.

The first contribution was the design and implementation of a two-electrode intermittent-harvesting interface for PMFCs. The proposed architecture connects the PMFCs to an energy-harvesting DC–DC converter while imposing the harvesting timing externally. This distinguishes it from capacitor-based and hysteretic approaches, where the harvesting and recovery intervals depend on the instantaneous cell state. By decoupling the imposed timing from the PMFC recovery dynamics, the interface enables frequency and duty cycle to be studied as controlled variables while preserving the practical relevance of converter-based two-electrode energy extraction. The same timing architecture can also support interleaved extraction from multiple PMFCs, providing a path toward scalable PMFC power management.

The developed interface was then used to evaluate whether intermittent operation can improve harvested power in a practical harvesting configuration. Preliminary timing exploration at 50% duty cycle indicated that seconds-scale periods were favorable for the tested cells, with the best results obtained around 5–10 s. In the 9-day experiment, reducing the intermittent period from 20 s to 10 s coincided with a strong increase in the average corrected PMFC-side power of both intermittent cells. Cell I increased from 47.2 μW to 122.4 μW, corresponding to a 159% increase, while Cell H increased from 19.4 μW to 67.6 μW, corresponding to a 248% increase. During the final interval, Cell I and Cell H exceeded the continuously loaded reference by 185% and 57%, respectively. Since the continuous reference also increased during Days 7–9, from 32.8 μW to 43.0 μW (+31%), the results support controlled intermittency as a promising strategy, but do not isolate the causal contribution of the period change from biological or environmental dynamics.

The second contribution was the use of EIS to observe how PMFC impedance evolves under different loading histories. The measurements were designed to be representative of harvesting operation rather than relaxed open-circuit operation: the anode and cathode polarization points used for three-electrode EIS were derived from the electrode potentials reached during two-electrode harvesting at 270 mV. In addition, the frequency sweep was bounded at low frequency to remain compatible with the energy-aware constraints identified in the previous battery-free EIS implementation. This links the laboratory EIS measurements performed in this chapter to the longer-term goal of embedded, self-powered PMFC state observation.

The impedance results showed that the loading profile leaves measurable signatures in the PMFC response. A common equivalent-circuit model could be fitted to both anode and cathode spectra, providing a consistent basis for comparing ohmic, interfacial, capacitive-like, and transport-related contributions across cells and operating conditions. The fitted parameters were interpreted comparatively rather than as unique physical descriptors, because the observed spectra reflect coupled effects of loading history, hydration, electrode state, transport, and biological activity. The analysis also showed that hydration events can strongly affect the impedance response, confirming that electrical loading and environmental state must be considered together when interpreting PMFC measurements.

Overall, the chapter shows that PMFC harvesting and PMFC state observation should be treated as a coupled problem. Controlled intermittent harvesting was associated with higher average extracted power under selected timing conditions, while EIS provides additional observables that reveal how the cell response changes under different loading histories. Together, these results provide preliminary evidence that comparable loading-history-dependent electrochemical changes also occur during practical two-electrode harvesting, as reflected in measurable electrode-level impedance signatures. At the same time, the experiments do not identify the biological origin of the observed impedance and power variations. Future work should therefore combine the proposed harvesting and EIS platform with longer experiments, larger PMFC populations, controlled hydration protocols, microbiological and plant-physiology measurements, and fully autonomous self-powered operation. This would enable PMFC-specific maximum power and energy tracking strategies in which operating voltage, intermittent timing, and diagnostic measurements are jointly optimized.

Chapter 4

CapDYN: Dynamic-Capacitance Architecture

Parts of this chapter are adapted from the author's previously published work on CapDYN [1] and from the related IEEE SENSORS work on zero-failure task execution in batteryless sensors [2].

Chapters 2 and 3 addressed the source side of the battery-free energy path by developing and experimentally evaluating controlled PMFC energy harvesting. This chapter moves to the storage and delivery side: once energy has been harvested from a weak and variable source, it must be buffered and made available to embedded workloads with heterogeneous energy requirements.

In conventional battery-free systems, harvested energy is accumulated in a fixed capacitor or supercapacitor, which supplies the microcontroller, sensors, and communication circuitry during active phases. Because ambient sources are often transient or sporadic, these nodes typically operate intermittently, alternating between charging and execution phases as the storage voltage moves between upper and lower operating thresholds [3, 6]. The choice of the storage capacitance is therefore a key design issue for intermittent systems [87].

A fixed capacitor creates a fundamental tradeoff, as the storage capacitor's size strongly influences both the device's active time

and the frequency of the power failures. A large capacitance can sustain energy-demanding operations and extend active time, but increases cold-start and recharge latency; a small capacitance improves responsiveness and mitigates long cold starts, but may not provide enough energy for longer sensing, processing, or communication tasks. Although a fixed capacitor can be sized for a specific operation, it is poorly matched to heterogeneous workloads, where different tasks require different energy budgets and response times. This mismatch can lead either to unnecessary charging delays for low-energy tasks or to insufficient stored energy for high-energy tasks.

For this reason, *variable-capacitance* storage architectures have been proposed to support workloads with different energy requirements while reducing unnecessary charging time. Two opposing capacitance-control approaches emerge from the literature. The first approach is to leverage *autonomous* control. Gogolou et al. [88] and El Mahboubi et al. [89] autonomously combine different sized capacitors in series or parallel to offer two capacitance levels, either a small capacitance for fast charging or a large capacitance for long on-time, respectively. The second approach is to rely on *software* control of capacitor size[21, 22, 90]: Yang et al. proposed Morphy [22]. Morphy is a software-defined configurable capacitor bank whose capacitors are connected via switches controlled by a *dedicated* microcontroller (MCU).

Capacitance steps are crucial for perfectly matching each task to its minimum capacity requirement, to minimize charging times. Energy efficiency is also paramount in battery-free ecosystems: the variable capacitance infrastructure must consume as little as possible to maximize the amount of harvested energy which is actually supplied to the load. Although software control enables the use of advanced control policies and grants a higher number of capacitance levels compared to the available autonomous solutions, this feat comes at the cost of an increased energy overhead.

Instead of having a variable capacitor bank, other works solve the limitations of a fixed capacitor when supplying different tasks by *federating* the energy storage [23, 91, 92, 24]. This consists in dimensioning a single fixed capacitor for each peripheral, providing control as to which capacitor to charge. Although federation can be an interesting solution, it still faces the same limitations of fixed capacitors at the peripheral level, if multiple tasks use the same

peripheral. Moreover, federation introduces the so-called *energy fragmentation* problem [24], as federation dilutes available energy across capacitors, so that even if the total stored energy would be sufficient for executing a task, operative voltage is not reached at the required peripherals. Solving this requires sub-optimal charge transfers among capacitors [24]. Because of this reasons, we focus on developing a centralized variable storage solution, rather than a federated one.

Problem statement. An ideal energy storage system should perfectly match the task at hand by providing just enough capacitance to sustain it, so that charging times are minimized. The capacitance control should minimally impact the energy efficiency of the variable storage. While the software control of state-of-the-art reconfigurable storages offers greater flexibility, it reduces the energy efficiency, as active MCUs are power-hungry components. On the other hand, the autonomous control can provide capacitance scaling with less power draw, especially if integrated. Besides, software intervention ties *efficiency* and *reactivity*: the more the system needs to be reactive to changes, the more the MCU needs to be active and the less energy efficient the variable storage is. Autonomous (i.e., not MCU-driven) capacitance scaling solves the reactivity-efficiency trade-off, as reactivity does not lead to an increased power consumption. This is essential to allow intermittent systems operating on µW-level transient power inputs to collect even the shortest bursts of incoming energy.

Challenges. For the purpose of explaining the challenges of building a variable capacitance system, consider a simple illustrative design - which we term D0 - similar to the one proposed in [21] that consists of a variable number of parallel capacitors. While in [21] parallel capacitors are set by a software policy, in the D0 autonomous design we start charging a single capacitor and *progressively scale up* the bank capacitance by adding a capacitor in parallel once a threshold charging voltage is reached, in order to make room for more energy. As soon as we store sufficient energy for a certain task, we start discharging the bank of parallel capacitors. This design introduces **two** important design challenges. **I)** When scale up is triggered, adding an initially discharged capacitor in parallel to a charged capacitor gives way to *energy waste* due to dissipative voltage equalization transients and potential *power failures*, as the overall voltage might stabilize below the minimum

supply level of the system. This can lead to undefined behavior of the load or even trigger a Brown Out Reset (BOR). **II)** Once the bank supplies the task and discharges to the minimum load supply voltage level, the energy below the voltage threshold can no longer be used to provide an operational voltage. This "leftover" energy cannot be consumed by the load, yet it must be harvested upon cold start to begin supplying the load, thus reducing the energy efficiency of the storage. With the aforementioned D0 design, we cannot tap into the leftover energy, which cannot be used to do useful work and unnecessarily burdens cold start.

Requirements. An ideal design for a reconfigurable capacitor bank should meet the following requirements: ⓪ It should provide granular capacitance steps in order to offer the perfect amount of storage for the next task in queue, to minimize charging times. ① It should minimize the leftover energy in the capacitor bank, once task execution is done. ② It must maintain the operating voltage during capacitance reconfigurations, without incurring energy waste and/or BOR. ③ It should perform the storage scaling at a low-energy cost to maximize the storage's *energy efficiency* (i.e. the ratio between output and input energy). ④ It must be *instantly reactive* to ambient power dynamics and load demands.

Chapter contributions. This chapter presents *CapDYN*, an ultra-low power energy storage design that *autonomously* adapts its capacitance to the declared energy consumption of the next task in queue. *CapDYN* overcomes the limitations of a fixed energy storage, which can only be tailored to the most energy consuming task within an application (typically, transmission), imposing long recharge times for smaller tasks (e.g. sensing, computing). CapDYN combines, within a single architecture, autonomous progressive capacitance scaling, task-matched storage stages, and a hardware energy-state interface for energy-admissible task execution. *CapDYN* begins charging at its minimum capacitance, to speed up cold start and allow conversion electronics to enter their most efficient conversion range as soon as possible. Each time the bank charges up to the upper voltage threshold, *CapDYN* then *scales* its capacitance up to the requested, task-specific level by seamlessly switching its bank of capacitors without the need of a dedicated MCU. Once the requested energy level is available, the task is launched and *CapDYN* gradually (i.e. each time the bank discharges to a lower threshold voltage) scales down its capacitance

back to its minimum level, maintaining the load's voltage range without output regulation. The load's range is maintained as voltage is boosted at each capacitance scale down step. Once the task is finished executing at the minimum capacitance level, leftover energy below the load's minimum operative voltage is minimized, resulting in significant improvements of the cold start times.

At the architectural level, *CapDYN* is designed to be readily extensible to additional stages: its control logic scales linearly, with a fixed incremental control-hardware cost per added stage. Practical scalability of the complete system, however, is constrained by the capacitor-bank topology. Of the two designs considered in this work, one exhibits linear component-count scaling, whereas the other scales exponentially with the number of stages, as detailed in Sec. 4.2.2. The resulting hardware implications, including leakage, parasitic losses, and implementation power, are experimentally characterized in Sec. 5.1.4. Moreover, additional stages are beneficial only when the resulting energy levels are suitably matched to the target workload, as investigated in the application-level evaluation of Sec. 5.2. Architectural extensibility therefore does not imply unbounded practical scalability: the evaluation in Chapter 5 characterizes the topology- and workload-dependent tradeoffs that determine a suitable stage count.

CapDYN can interface with harvesting sources whose output can be appropriately conditioned by source-specific conversion electronics. The upstream front end must rectify the harvested signal when required and buck or boost its voltage to provide a DC input capable of charging the main bank up to the upper switching threshold V_{SH}. This requirement follows from *CapDYN*'s direct connection between the storage bank and the load, without a downstream output boost stage. Under this assumption, source compatibility is determined by the availability of a suitable conditioning front end rather than by the harvester type itself.

Additionally, we develop a lightweight intermittent-computing runtime that works in conjunction with *CapDYN*. The runtime receives energy-aware signals from *CapDYN* and maps the energy consumption of tasks to the energy available at each capacitance level of the bank, enabling *energy-admissible task execution*: tasks are released only when the current storage state can sustain their completion. In short, the main contributions of this chapter are as follows:

- We present *CapDYN*, a dynamic energy storage solution which enables *autonomous* scale up/down of its overall capacitance level based on its state of charge, without needing any external MCU.
- We provide a *CapDYN* lightweight runtime that exploits the bank's energy state interface to eliminate load-side sleep-mode consumption during inter-task charging by power-gating the load until the energy required by the next task is available, and to efficiently and intermittently execute energy-aware, task-based programs.
- We release the *CapDYN* hardware design, LTspice model, and software runtime as open-source to support reproducibility and foster the adoption of batteryless energy-harvesting systems [93].

The remainder of the chapter is organized as follows. Sec. 4.1 reviews prior work on intermittent systems and dynamic capacitance and motivates the design requirements for dynamic energy storage. Sec. 4.2 presents the *CapDYN* hardware architecture, including the progressive capacitance-scaling mechanism, capacitor-bank design, power-flow management, and energy-state interface. Sec. 4.3 introduces the energy-aware runtime and its task-based execution model, while Sec. 4.4 describes the hardware, LTspice, and software implementations. The chapter concludes by defining the task-to-stage mapping used to dimension *CapDYN* for heterogeneous workloads. The experimental and system-level evaluation of the resulting architecture is presented in Chapter 5.

4.1 Dynamic capacitance: related works and motivation

Most battery-free energy harvesting devices feature a single, fixed-size centralized capacitor, which is generally as large as the most energy-consuming operation that must be executed without power failure [95, 96]. This introduces a delicate balance if the storage must support operations with varied levels of energy consumption: while a larger energy reservoir can support more substantial tasks, it concurrently prolongs the recharge time for shorter tasks, thereby diminishing the system's overall responsiveness [21].

Table 4.1: Comparison between *CapDYN* and representative prior work in storage architectures. Adapted and extended from [1].

	Dynamic Charging Time	Reduced Leftover Energy	Task-Storage Mapping	Energy-Admissible Release	Auto-Reconfigurability
Single Capacitor	✗	✗	✗	✗	✗
Capybara [21]	✓	✗	✓	✓	✗
Reconf. SC [94, 88], Variable Storage [89]	✓	✗	✗	✗	✓
UFoP [23], Flicker [91]	✓	✗	✗	✗	✗
Stash [92]	✗	✗	✗	✗	✓
Multi-storage EMU [90]	✓	✗	✓	✓	✗
Morphy [22]	✓	✓	✓	✗	✗
CapDYN	✓	✓	✓	✓	✓

As highlighted in prior articles [22, 21, 30, 97], a batteryless system faces multiple power supply-related challenges. These include achieving operational voltage from an empty storage (cold start), sustaining tasks with different energy requirements, using stored energy efficiently, and managing intermittent execution. Recent works have focused on addressing these challenges by providing variable storage upon software request [21, 22] or autonomously [94, 88] switching between a small and a large capacitor to achieve faster startup [89], or by considering per-peripheral capacitors [23]. Among these works, reconfigurable capacitor architectures [22, 21] provide a clever way to solve the fast recharge/sustained operation tradeoff. However, connecting multiple capacitors and enabling their software-based reconfiguration introduces additional overhead, increasing complexity and power requirements.

4.1.1 Differences from the State-of-the-Art

Tab. 4.1 summarizes the main differences between *CapDYN* and prior reconfigurable-storage architectures. Beyond combining features previously explored in isolation, *CapDYN* introduces autonomous reconfigurability and exposes the storage state as a hardware-level execution contract, enabling both efficient energy management and energy-admissible task execution.

① Dynamic Charging Time and reduced leftover energy.

A simple solution to provide dynamic charging time is to provide two capacitance steps [94, 88, 89], one for fast startup, one for storing extra energy if the input power is greater than the end-node's active power consumption. While this approach is an important example of autonomous capacitance scaling, the purpose of providing a larger capacitance step is to store excess energy without burdening cold start, not to support larger tasks. Indeed, larger tasks rely-

ing on the larger capacitance steps might never execute, unless the power input is higher than the node's active power consumption.

This shortcoming is improved with Capybara [21], where storage size is explicitly matched to task energy request, be it large or small, and the charging of higher capacitance levels is decoupled from active load power consumption, as the load is only activated when sufficient energy for task execution has been harvested. When scheduling a task, Capybara connects a certain number of parallel capacitors sufficient to match the task's energy requirements (termed *energy mode*). The storage is only modified when scheduling a task, then it acts as a fixed-size capacitor (i.e. it does not reconfigure throughout charging and task execution). Given this reconfiguration scheme, to minimize leftover energy, Capybara must employ an output voltage boost to generate an operative voltage as the bank dips below the load's lower operative voltage. This approach surely favours simplicity, but it comes at the cost of some extra infrastructural power consumption and conversion losses due to the output boost converter.

The multi-storage EMU of Mileiko et al. [90] can be viewed as a variation of Capybara's configurable storage approach. Capybara maps each task to a task-specific set of capacitors connected in parallel, whereas Mileiko et al. map it to one selected capacitance C_s among a set of mounted storage options and to an operating voltage V_{op}. As in Capybara, this task-specific storage configuration is selected before execution and then remains fixed while energy is accumulated and consumed, rather than being autonomously reconfigured throughout charge and discharge. A programmable voltage monitor determines when the selected capacitor reaches V_{op}, while a DC–DC converter regulates the voltage supplied to the load. The combined capacitance-and-voltage energy representation provides finer control over the energy allocated to each task. However, only one storage capacitor is selected at a time, and no mechanism is described for recovering the energy remaining below the minimum system operating voltage or for reusing residual charge in a storage capacitor after it is deselected. Leftover-energy management is therefore not explicitly addressed.

Yang et al. further expanded the software-defined capacitance of Capybara through *progressive scaling* in Morphy [22], an approach that also distinguishes Morphy from the later task-selected storage architecture of Mileiko et al. The same idea, with an improved im-

plementation, is also employed in *CapDYN*. The progressive scaling mechanism consists in gradually reconfiguring the capacitor bank to scale up/down the bank capacitance level throughout charging/discharging. Charging begins at a minimum capacitance level, so that operative voltage is achieved as soon as possible. Then, capacitance is scaled up either when an upper charging voltage is reached (*CapDYN*) or when the bank is deemed optimally charged (Morphy), to make room for more energy. When energy is deemed sufficient for task execution, bank discharge begins, and capacitance is gradually scaled back down to the minimum level. This provides three advantages over task-selected storage approaches such as Capybara and Mileiko et al., which set up the storage configuration before charging/discharging and do not reconfigure it until a new task is scheduled. First off, operative voltage is achieved earlier, and then it is maintained through progressive scaling. Conditioning electronics circuits have a much higher efficiency (5% vs 90% for commercial models such as Texas Instruments' BQ25505) when their output voltage - which is where the dynamic capacitor is placed - grows above a cold start voltage threshold. Progressive scaling allows *CapDYN* and Morphy to overcome cold start much sooner than a fixed capacitor tailored for the same task, meaning that conversion electronics work more efficiently sooner, resulting in more stored energy. Secondly, at each scale down, capacitors within the bank are reconfigured to boost the bank's voltage, until the minimum capacitance is reached: leftover energy - energy below operative voltage - is thus effectively reduced with respect to the fixed capacitor needed for an equivalent task execution. Leftover energy below the load operating voltage cannot be used for useful work but must be re-harvested at the next cold start; reducing this residual energy therefore improves both energy utilization and restart time [1]. Thirdly, voltage boosting is achieved without using an output boost component as in Capybara, which introduces additional consumption and inefficiencies.

CapDYN further improves the efficiency of the progressive scaling mechanism with respect to Morphy. Morphy implements progressive scaling by charging its entire bank, and then separating a subset of task-tailored capacitors during execution. Charge-pump boosting is then employed between the core array and the task capacitors, to maintain the task capacitors in the load's voltage range. While this improves the flexibility of Morphy in terms of load volt-

age range and scaling policies, it also causes large dissipative transients due to frequent charge transfers among differently charged capacitors, which largely limit the efficiency of the boosting process. *CapDYN* does not feature a core/task capacitor separation: instead, *CapDYN* avoids the inefficiencies they carry by tailoring its whole bank's capacitance levels to match the different task energy requests. This further speeds up the charging process with respect to Morphy, where overall bank capacitance is not explicitly matched to a single task, as the task-matching job is left to the core/task array separation. As a bonus, by matching capacitance steps with task requests, we obtain a hardware *energy state interface*: by autonomously monitoring the digital lines used to pilot the bank switches, *CapDYN* can fully power-gate its load until it reaches the stage needed for the next task's execution. This energy state interface is thoroughly explained in Sec. 4.2.3.

② **Auto Reconfigurability.** Software-controlled dynamic bank solutions allow greater scheduling flexibility, but do not solve the reactivity-efficiency trade-off. Indeed, if reconfiguration is tied to the activation of a power-hungry MCU [21] or periodic state of charge measurements [22], the more the system will require reconfiguration, the more energy will be wasted. Autonomous control reduces power consumption, but also solves the reactivity-efficiency tradeoff: active monitoring and reconfiguration do not come at an added cost, unlike with software-controlled solutions, as evaluated in Sec. 5.1.5. For these reasons, *CapDYN* was designed to operate autonomously, without an MCU. Although the current prototype is implemented using discrete components, such as voltage comparators and logic gates, *CapDYN* could be integrated into a single chip, further reducing power consumption and improving its energy efficiency. Among the architectures reviewed in this section, *CapDYN* is distinguished by the combination of auto-reconfigurability, scheduling support, and a control architecture that is linearly extensible with the number of stages.

Other autonomous solutions presented in literature are focused on harvesting excess energy, rather than providing dynamic capacitance. In STASH [92], the front-end, federated capacitors are tailored to sustain peripherals throughout execution, while a back-end capacitor is available to autonomously collect energy when the energetic needs of the load are satisfied (i.e. the front-end storage is full and task execution needs to wait). This mechanism, although

inefficient due to the charge transfer through a Shottky diode from the back-end capacitors to the front-end capacitor, manages to store extra energy without slowing down cold start, as the back-end capacitor is connected to the harvester output only in times of energy abundance. A similar result to Stash's is obtained in [94, 88, 89], but extra energy is only collected if the power input is higher than the load's active power consumption. *CapDYN*'s autonomous control, coupled with its *progressive scaling*, manages to reactively expand, even if the next task energy level has been reached and task execution needs to wait, essentially providing Stash's functionality without extra hardware.

③ **Hardware support for Energy Awareness** Energy-aware execution requires knowledge of both task energy demand and the storage state of charge. While the former can be obtained through offline profiling, the latter requires active measurements, threshold monitoring, or dedicated hardware support. UFoP [23] and Morphy [22] actively poll storage voltage, incurring considerable energy overhead due to MCU intervention, whereas Flicker [91] avoids this overhead through interrupt-based monitoring. Mileiko et al. [90] similarly avoid periodic MCU polling through dedicated threshold-monitoring hardware. Flicker couples energy awareness and charge prioritization through two programmable voltage monitors per peripheral: a charge controller and an interrupt controller. The charge controller enables charging of a peripheral capacitor when the first-stage capacitor, which powers the core, rises above a programmable threshold; lower thresholds correspond to higher peripheral priority. The interrupt controller signals when the peripheral capacitor has reached the voltage required for the corresponding task to execute. By reconfiguring these thresholds, Flicker supports energy-aware scheduling and priority assignment. However, this flexibility comes at the cost of replicating monitoring hardware across peripherals and introducing charge-prioritization complexity. Moreover, the lack of implementation detail on charge transfer across capacitors makes it difficult to assess the efficiency and scalability of this approach, especially when peripheral voltage ranges interact [92]. Mileiko et al. [90] also leverage a programmable voltage monitor, but unlike Flicker's per-peripheral monitoring and charge-prioritization scheme, they directly monitor the selected centralized storage against the task-specific voltage threshold required for the next task's execution.

In contrast to Flicker's and Mileiko's programmable-threshold approaches, *CapDYN* does not require task-specific voltage-threshold configuration for energy awareness. Instead, the *CapDYN* bank's digital lines serve a dual purpose: they drive bank reconfiguration and expose the current capacitance step—which also represents the storage state of charge—through the *energy state interface* (Sec. 4.2.3). Through hardware/software co-design, each capacitance step is associated with the set of tasks it can sustain. As a result, the runtime obtains immediate feedback on which task can execute by simply observing the current storage stage.

Thanks to this *energy state interface*, the load running *CapDYN*'s runtime (Sec. 4.3.3) is activated only when the main bank stores enough energy to complete the next task. This greatly improves efficiency by removing the overhead of sleep-mode power consumption and periodic polling of the bank's state of charge, both of which burden software-controlled dynamic capacitors. In addition, *CapDYN* centralizes charge in a single bank, avoiding the need to multiplex energy across federated capacitors. This removes the complexity and overhead of hardware charge-prioritization schemes such as Flicker's.

④ **Time-constrained execution.** Batteryless devices need to execute tasks which pose constraints not only on *how much* energy they need to successfully execute, but also on *when* they need the requested energy. Such *time-constrained* tasks need to execute *reactively*, as an immediate response to an event. *Recharge time* determines whether a time-constrained task will execute reactively, as the system is unresponsive during recharge [98, 99]. As discussed in the introduction of this chapter, a fixed buffer cannot reactively respond to tasks with heterogeneous energy requirements: it either imposes long charging times for shorter tasks or provides insufficient energy for longer ones. While energy-constrained tasks only require visibility into the storage state of charge, time-constrained tasks additionally require either *reduced recharge time*, so that the buffer reaches operating voltage as soon as the requested energy is available, or *pre-charging*, so that energy for reactive execution is reserved in advance. Catnap [97] logically partitions a fixed-size energy buffer, reserving charge for reactive responses while executing an energy-constrained task queue when extra energy is available. This allows fine-grained control over the pre-charge level, but requires continuous energy monitoring, which is inefficient. Similarly,

Capybara [21] supports hardware pre-charging for time-critical responses by charging a section of its bank to the required level and then isolating it until the reactive task arrives. However, the isolated bank may gradually lose charge due to leakage. *CapDYN* supports both pre-charging and recharge-time minimization. Pre-charging is achieved by preserving the lower capacitance steps for reactive responses, with energy reservation enforced through *CapDYN*'s digital output lines. Unlike Capybara, the reserved energy is not isolated from the main bank, but remains in the central storage, which continues to charge and therefore counteracts leakage. Recharge times are further reduced through *CapDYN*'s progressive scaling mechanism, which matches the capacitance step to the reactive task demand while also improving energy efficiency and first completion advance (Sec. 5.2). Moreover, if persistent timekeepers such as CHRT [100] are integrated into *CapDYN*, recharge time can also be exploited to infer average harvested power, opening the possibility of adapting the schedule to the input-power conditions.

⑤ Energy-admissible Task Execution and Scheduling Support. Intermittent computing is fundamentally an *energy-awareness* problem. Under harvested and unstable power, execution correctness and efficiency depend not only on preserving state across outages, but also on deciding whether the energy currently stored in the buffer is sufficient to safely start, continue, or defer computation. Existing approaches differ precisely in *how* this energy awareness is obtained and *how* it is used: to trigger checkpointing, to tolerate failed task executions, or to predict safe task release.

Traditionally, checkpointing [101], which saves the system state to non-volatile memory (NVM) during energy depletion and restores it upon power recovery, has been the dominant approach. Early implementations rely on static checkpoints embedded in the source code. For instance, [102] proposes a system that analyzes application code to optimize checkpoint placement. However, static checkpointing often leads to energy waste when checkpoints are bypassed without being used, or to failures when power drops before a checkpoint is reached. Hibernus++ [103] mitigates these limitations through adaptive checkpointing, which introduces the use of energy awareness as a mechanism to decide when execution is no longer safe and state saving must be triggered. Energy awareness is achieved by capacitor-voltage monitoring through an ADC and by energy-margin estimation to determine when hibernation

is necessary, rather than checkpointing at a fixed threshold as in [7]. Subsequent works such as Chinchilla [104] refined this approach by introducing predictive checkpointing mechanisms that leverage energy and execution models to anticipate future energy availability and schedule hibernation decisions accordingly. To further reduce overhead, selective checkpointing schemes [105] target only live data, while Dice [106] saves only modified memory regions. Nevertheless, checkpointing-based systems still incur non-negligible time and energy overheads, which limit their effectiveness under highly intermittent power.

Task-based execution models instead shift the focus away from energy-aware state preservation and toward *failure-tolerant execution*. They decompose applications into atomic units that can be safely re-executed after reboot, as tasks are designed so that they can be safely re-executed from the beginning without causing incorrect or inconsistent updates to non-volatile memory, thereby avoiding explicit state saving during normal execution. Chain [107], InK [108], and Alpaca [109] follow this approach through specialized task abstractions and communication channels that preserve consistency across power cycles. LATICS [110] and Coala [111] improve upon static tasking by adapting task boundaries using indirect energy-awareness signals: recent execution outcomes in LATICS and elapsed execution time in Coala. These mechanisms reduce overhead and mitigate non-termination, but they still provide limited energy awareness: tasks are launched without guaranteed completion, so repeated reboots and energy waste remain possible when input power is insufficient.

Checkpointing-based and task-based approaches represent the closest intermittent-computing reference classes for the proposed solution. They address the same general problem as *CapDYN*, namely forward progress under intermittent power, but they act mainly at the software execution and recovery level. Checkpointing preserves program state across power failures, whereas task-based models structure applications into atomic units to simplify recovery and consistency. In contrast, *CapDYN* addresses the complementary storage and activation problem: it exposes discrete energy states and enables execution only when a profiled unit of work is energy-admissible. The proposed system should therefore not be interpreted as a replacement for checkpointing or task-based programming models. Rather, it provides storage-level support that can be

combined with them, allowing such software mechanisms to benefit from reconfigurable storage and energy-admissible activation.

Predictive and optimization-based approaches make energy awareness explicit again, but use it for *task admission* rather than only for state preservation. Their goal is to decide whether a task should be started at all, based on measured, estimated, or predicted energy availability. Karimi and Kim [112] propose model-based prediction of the stored energy evolution over time so that the node is woken up only when sufficient energy has been accumulated. ETAP [113] estimates available energy using hardware-level measurements such as capacitor voltage and compares it against compile-time-derived probabilistic energy bounds, enabling tasks to be scheduled only when sufficient energy is available with high confidence. Delgado and Famaey [114] formulate batteryless task scheduling as a Mixed Integer Linear Program that uses task timing and energy parameters together with predicted harvested energy over a look-ahead window to decide task execution and node duty-cycling. Their formulation yields a globally optimal schedule only under full knowledge of future tasks and harvested power over the considered horizon, and is therefore best understood as an offline upper bound and a basis for practical heuristic schedulers. More generally, predictive and optimization-based approaches reduce task failures at the cost of relying on prediction accuracy and potentially complex models, which have limited applicability in resource-constrained environments.

In contrast to the previously introduced software-side approaches discussed above, *CapDYN* treats energy awareness not as a software-side estimate used to trigger checkpointing, tolerate failures, or predict safe execution opportunities, but as a *hardware-enforced execution contract*. *CapDYN* shifts intermittent execution away from checkpointing, speculative task retries, and prediction-heavy scheduling toward *energy-admissible* task release enforced by the storage/runtime interface itself. *CapDYN* couples task execution with a dynamically reconfigurable centralized capacitor bank through its *energy state interface* (Sec. 4.2.3). Through offline task-to-stage mapping, each task is associated with the minimum capacitance stage that is expected to guarantee completion under the profiled energy bound and safety margin. At runtime, the node determines whether the next task is energy-admissible simply by observing the current energy state exposed by the bank, thus eliminating the need for explicit online energy estimation, voltage polling, or predictive

scheduling. Energy awareness is therefore provided implicitly by the storage itself, which substantially reduces both runtime complexity and monitoring overhead.

Building on the hardware support described in Secs. 4.2.3 and 4.3.3, this mechanism also strengthens support for both energy-constrained and time-constrained tasks. As discussed in ④, *CapDYN* can reserve energy for reactive responses by preserving lower capacitance steps and can reduce recharge latency through progressive scaling, thereby addressing both *how much* energy a task requires and *when* it must be delivered. Unlike Flicker, this does not require per-peripheral charge prioritization; unlike Capybara, reserved energy remains in the main bank rather than in an isolated storage partition.

Mileiko et al. [90] provide a hardware-supported task-admission model closely related to that of *CapDYN*, associating profiled tasks with task-specific energy states and power-gating the transient load until the requested state is available. A key distinction lies in how this requested energy state is retained across power failures. Mileiko et al. integrate a dedicated State Retention Unit (SRU), in which the selected energy state (C_s–V_{op}) is maintained in volatile latches during normal operation and stored in NVM for autonomous restoration after an outage. In *CapDYN*, the requested stage is instead held in a volatile wake-up register; after complete control-power loss, this register returns to a user-configurable default stage, from which the load can restart, recover its application state from MCU non-volatile memory, and program the required stage again. Mileiko et al. therefore directly retain the exact energy-management state, whereas *CapDYN* avoids dedicated control-side NVM and restoration circuitry and reconstructs the requested state after restart.

The choice between direct non-volatile retention and restart-time reconstruction introduces a tradeoff between power-failure recovery, control complexity, and continuous power consumption. Mileiko et al.'s non-volatile restoration mechanism requires only 0.31–0.4 μJ for state recovery and 0.07–0.1 mJ for state update. However, the SRU alone draws approximately 3 μA during normal operation, while the complete EMU includes additional control and power-management circuitry whose total quiescent consumption is not reported. A direct quantitative comparison with *CapDYN*, whose complete control infrastructure consumes 4.15–9.07 μW, is therefore not possible. Mileiko et al. thus obtain direct non-volatile recov-

ery of the energy-management state through additional retention circuitry, whereas *CapDYN* favors lower control complexity and reconstructs the requested stage after a complete control-power loss: as detailed in Sec. 4.2.3, power-on reset first restores a configurable default wake-up stage, after which the restarted load recovers its application state from MCU non-volatile memory and programs the stage required by the next task.

By enforcing energy-admissible execution at the storage/runtime boundary, *CapDYN* overcomes the uncertainties of predictive scheduling and virtually eliminates the overhead of checkpointing. Under normal intermittent operation, execution suspensions due to insufficient stored energy are confined to inter-task charging intervals and coordinated through the `EOT` signal of the energy state interface. Consequently, once started, tasks are guaranteed to complete, and the system only needs to preserve minimal inter-task state to ensure forward progress. In this sense, *CapDYN* turns energy awareness from a software-side estimation problem into a hardware-enforced execution contract between the storage and the runtime.

4.1.2 Centralized and federated solutions

A radically different approach to centralized, variable capacitance is to *federate* the energy storage across peripherals, by dedicating a fixed-size capacitor to each peripheral. By only activating the charging of the required peripheral capacitors for a given task execution, overall storage capacitance can be reduced with respect to a fixed-size centralized capacitor, especially in the case of applications with diverse energy requirements. This approach has three major shortcomings with respect to centralized solutions. Firstly, even though the time to operative voltage is minimized across all required peripherals, task execution needs to wait for the charging of the slowest peripheral. This means that the peripherals which already have achieved operating voltage sit idle and draw power from their storages. With a centralized solution, by tailoring the capacitance to a task, we minimize not only the time to operative voltage, but also the time to task execution, so when useful work is done, without having to consider peripheral charging synchronization delays and energy waste. Secondly, in federated approaches energy *is fragmented* across peripheral capacitors, as debated in [24], which

means that even though the overall stored energy would be sufficient to launch a task, each peripheral is below operating voltage and cannot operate. Solving this problem requires sub-optimal charge transfers among different peripheral capacitors, which reduces the energy efficiency of the system [24]. Centralized storages such as *CapDYN* avoid fragmentation by design and do not require lossy defragmentation techniques. Besides for the purpose of defragmentation, the routing of charge from one capacitor to a differently charged one regularly takes place in federated storages (e.g. stash capacitor to peripheral capacitors [92], MCU capacitor to peripheral capacitor [23, 91]), and gives way to large energy transfer losses. Thirdly, federated storages dimension the peripheral capacitor according to the consumption of a single task using that peripheral. If a peripheral requires different energy budgets for different tasks, this approach inherits the shortcomings of fixed capacitors when scheduling tasks with varied energy draws. In the case of centralized systems such as *CapDYN*, Capybara and Morphy, we can remap different tasks using the same peripheral to different capacitance levels, overcoming the mentioned limitation of per-peripheral capacitors. Hester et al. [23] argue that, unlike federated storages, centralized solutions are vulnerable to the so-called "tragedy of the coulombs", which is what happens if a faulty task or component can access and wrongly deplete a shared storage. Centralized systems can defend themselves from this problem by implementing *energy watchdogs* [22], which power-gate peripherals through switches if the energy draw goes above the expected level. *CapDYN* is immune to this problem, as tasks are triggered as soon as the bank has sufficient energy to complete them. So, a bugged task will, at worst, not complete. Blacklisting can also be implemented to avoid re-execution. If pre-charging is activated, we can safeguard the pre-charged energy through an energy watchdog which leverages *CapDYN*'s capacitance step knowledge.

4.1.3 Other relevant works

Two recent works, Culpeo [115] and CapOS [116], propose interesting solutions which could be integrated with *CapDYN*. Culpeo [115] considers the voltage drop caused by the equivalent series resistance in capacitors. This voltage drop might lead to an unexpected power failure even though the stored energy remains in the capacitor.

CapOS [116] considers the capacitor degradation due to repeated charging and discharging operations. The benefits of the proposed solutions in these works can be applied to *CapDYN* to further improve its performance. We leave these aspects out of the scope of this thesis.

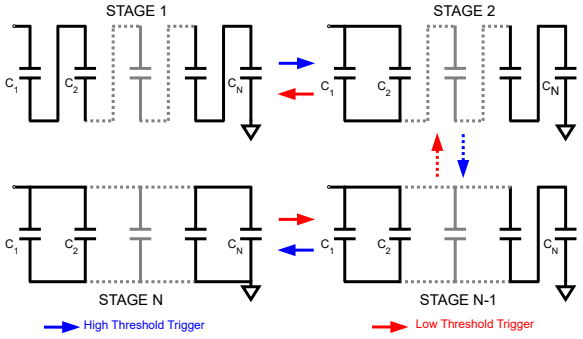
4.2 *CapDYN* hardware design

With *CapDYN*, we introduce a novel dynamic energy storage which *automatically* adjusts its overall capacitance based on its current state of charge. *CapDYN* executes these capacitance adjustment in steps though what we term *progressive scaling*, as shown in Fig. 4.1: upon system cold start, *CapDYN* begins charging at its minimum capacitance level, then it *progressively* increases its capacitance level each time an upper voltage threshold is reached. When the capacitance level required to execute the next task *without power failure* is fully charged (Sec. 4.2.3), *CapDYN* enables task execution: as the bank discharges, *CapDYN* decreases its capacitance level each time a lower voltage threshold is crossed, until the minimum capacitance level is discharged down to the lower voltage threshold. At this point, *CapDYN* has extracted all viable energy from its bank, having minimized leftover energy through the aforementioned progressive scaling mechanism. Through the scaling process, the output voltage is limited, as the output voltage is maintained between an upper and a lower voltage threshold. Besides this, the progressive scaling approach gives *CapDYN* the following advantages with respect to a fixed capacitor solution:

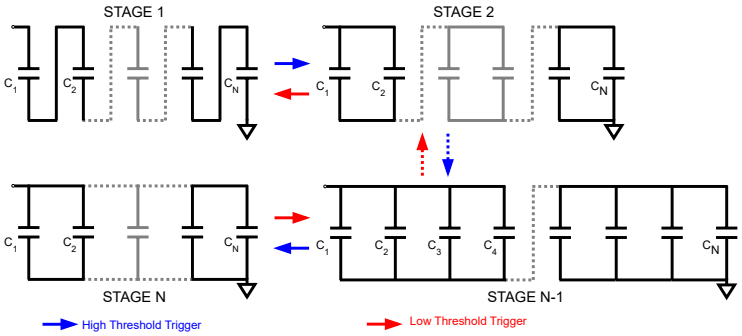
- By tailoring *CapDYN*'s capacitance steps to each task's consumption, we minimize charging times across tasks with varied energy requirements, all the while sustaining their execution without incurring power failures.

- *CapDYN* reaches operative voltage sooner and maintains it throughout charging, thanks to its *progressive scaling* approach: this means conditioning electronics operate in their most efficient range (that is, out of cold start) for more time.

- *CapDYN* can discharge the bank down to its minimum capacitance level, effectively reducing the leftover energy - that



(a) Design 1.



(b) Design 2.

Figure 4.1: The *progressive scaling* capacitor reconfiguration pattern for *CapDYN*, in two bank design variants. The initial configuration of both banks is *stage 1*, where all the capacitors are connected in series. As the stored energy increases, each design differently reconfigures the connections between the capacitors in successive steps to connect them in parallel. At *stage N*, all the capacitors of both banks are connected in parallel. Reproduced from [1].

is, energy that cannot be used for powering the load and thus unnecessarily burdens cold start.

CapDYN system's architecture is composed of four parts, which are discussed in the following section and represented in Fig. 4.2. **1)** The energy storage *scaling logic*, which enforces *progressive capacitance scaling*; **2)** the capacitor bank (or *main bank*) with its reconfigurable switches, piloted by the scaling logic; **3)** the *I/O flow control*, which orchestrates the power flow into and out of the main bank and **4)** the *energy state interface*, which allows the load to indicate at which energy level it wants to be powered, in order to execute its next task. The scaling logic is in charge of monitoring the main bank voltage and of accordingly piloting the progressive scaling of the bank's capacitance through the bank's switches. The proposed scaling logic is not tied to a single capacitor-bank topology and can be matched with different reconfigurable bank designs. In this work, as a proof of concept, we focus on two convenient capacitor-bank designs. The I/O flow control ensures that the main bank is only ever charged/discharged when the backup capacitor can sustain the scaling operations. The energy state interface allows the bank's power output to be activated by the I/O flow control only once the load's next task energy level has been achieved, in order to avoid wasting energy in load sleep mode power consumption. The power output is then disabled when the main bank is fully discharged, in order to avoid successive cold starts. Both the scaling logic and the I/O flow control are fully hardware-based, as they forego active control (i.e., an MCU) to carry out their tasks. As we will see in Sec. 5.1.5, this solves the reactivity-efficiency tradeoff mentioned in this chapter's introduction. Besides being flexible in terms of capacitor-bank design, *CapDYN* can interface with harvesting sources whose output can be conditioned by source-specific conversion electronics to provide a DC input capable of charging the main bank up to at least V_{SH}, through rectification and/or buck/boost conversion as required.

4.2.1 The progressive capacitance scaling logic

CapDYN's scaling logic relies on monitoring the main bank's voltage to estimate the stored energy within the bank and to autonomously trigger scale up/down of the bank capacitance. Capacitance scale-up/down is achieved by modifying the connections between the

bank's capacitors through a series of switches, which are controlled by the scaling logic. Capacitance scale-up triggers when the main bank's voltage grows above an upper threshold V_{SH}, while scale-down triggers when the voltage falls below a lower threshold V_{SL}. Note that $V_{SH} > V_{SL}$. In this way, the main bank is limited in the $[V_{SL}, V_{SH}]$ range at all times during normal operation. In its most basic form, the scaling logic is a shift register which shifts its bits at each scaling event. Each bit pilots a subset of the bank's switches to implement different overall capacitance steps. By using different bit combinations to control the bank's digital lines, *CapDYN*'s scaling logic can be adapted to different reconfigurable capacitor-bank designs whose stage transitions can be encoded through the available control lines.

Avoiding Brown Out, Overcharging and Loops in Progressive Scaling

Switching thresholds V_{SL}, V_{SH} require careful design in relation to the load's voltage range, V_L, V_H and to the biggest voltage drops/spikes associated to capacitance variations. Voltage spikes/drops upon transitions depend on the selected bank design, specifically on the reconfiguration strategy and on how the capacitors within the bank are dimensioned, as discussed in Sec. 4.2.2.

In order to directly (i.e. without direct output regulation) supply an operating voltage to the load throughout the progressive scaling operations, the main bank's output voltage must be limited to the load's operating range. Indeed, if capacitance scale-up/down gives way to a bank voltage spike/drop that exits the load's range, the load could be either damaged or suffer from a power failure. At the same time, if these voltage spikes are larger than the $[V_{SL}, V_{SH}]$ range, a *loop condition* occurs. In fact, if upon reaching V_{SH}, capacitance scale-up generates a voltage drop that reaches V_{SL}, scale-down is triggered right after scale-up, and scale-down comes with a voltage spike which newly exceeds V_{SH}, resulting in a new scale-up. In order to respect the load's voltage range throughout scaling and make as much room as possible to accommodate scaling-associated voltage spikes/drops, we select $V_{SL} = V_L$, $V_{SH} = V_H$. In the next paragraph, we highlight two capacitor dimensioning solutions which can be matched to different load voltage ranges.

4.2.2 Main bank design: switching strategy and capacitor dimensioning solution

CapDYN's capacitor bank is composed of multiple capacitors, whose connections are reconfigured through switches piloted by the scaling logic. A capacitor bank is fully defined by its *switching strategy* and by its *capacitor dimensioning solution*. The switching strategy defines how the capacitors of the bank are arranged in order to achieve varying bank capacitance steps, while the capacitor dimensioning solution expresses how each capacitor within the bank is dimensioned with respect to each other.

CapDYN is compatible with multiple bank designs. Fig. 4.1 shows the switching strategies of the two banks we evaluate in this paper. Design 1 (D1) ensures a large capacitance multiplication range with limited components, with a component count that increases linearly with the number of stages. Design 2 (D2) instead requires a twofold increase in component count for each additional stage, resulting in exponential component-count scaling, but features more energy-efficient transitions. The hardware-level differences between the two topologies are evaluated in Sec. 5.1.4, while their application-level impact is analyzed in Sec. 5.2.

Switching strategy: D1's *switching strategy* (Fig. 4.1a) features all bank capacitors connected in series in the initial stage/level, setting the minimum capacitance to enable quick cold-start overcoming. As the bank charges, its capacitance is progressively increased: each time the voltage reaches V_{SH}, one capacitor is connected in parallel with C_1. The maximum capacitance level is reached when all capacitors are connected in parallel to C_1. Bank discharge (task execution) is enabled once a target energy level—corresponding to a specific capacitance level—is reached (Sec. 4.2.3). During discharge, the process is reversed: each time the voltage falls to V_{SL}, one capacitor is removed from the parallel configuration and reconnected in series, providing the benefits discussed in Sec. 4.1.1.

D2's *switching strategy* (Fig. 4.1b) follows a similar initial condition, with all capacitors connected in series. However, its scaling mechanism differs significantly. When the bank voltage reaches V_{SH}, capacitors are reconfigured in pairs, forming parallel groups of two capacitors, connected in series with other such groups. Upon subsequent crossings of V_{SH}, these groups connected in series are recursively merged: pairs of two-capacitor groups are combined into

four-capacitor parallel groups, and so on.

Eventually, the maximum capacitance is achieved when all capacitors are connected in parallel. During discharge, the process is inverted: each time the voltage reaches V_{SL}, each parallel group is split into two and the resulting subgroups are reconnected in series, ultimately returning to the initial fully series configuration.

From this description, it is evident that D1 scales its bank capacitance linearly with the number of stages, whereas D2 exhibits exponential scaling, as each additional stage requires a twofold increase in capacitors. As discussed in Sec. 5.1.4, this difference has important implications in terms of both power consumption and cost.

Capacitor-bank dimensioning and stage-energy availability: For an N-stage D1 *CapDYN*, the capacitor bank is dimensioned using a binary-weighted array, where all capacitors are expressed as multiples of a *unit capacitance*, C_1. Specifically,

$$\begin{aligned} C_2 &= C_1 \\ C_i &= 2^{i-2} \cdot C_1, \quad i = [3, N] \end{aligned} \tag{4.1}$$

At stage j, the bank can be interpreted as the series connection of two sub-networks, a parallel block composed of the first j capacitors, and a series string composed of the remaining $N - j$ capacitors. Under this dimensioning, the equivalent capacitance of the bank at stage j can be expressed in closed form as

$$C_{\text{eq}}^{\text{D1}}(j; N) = \frac{C_1\, 2^{j-1}}{3 - 2^{\,j+1-N}}, \qquad (1 \le j \le N), \tag{4.2}$$

In the D2 bank (Fig. 4.1b), all capacitors have equal size C_1. At stage j, capacitors are grouped into parallel blocks of size 2^{j-1}, which are then connected in series. The resulting equivalent capacitance can be written in closed form as

$$C_{\text{eq}}^{\text{D2}}(j; N) = C_1\, 2^{2j-N-1}, \qquad (1 \le j \le N), \tag{4.3}$$

which increases exponentially with the stage index until all capacitors are connected in parallel at $j = N$.

Thus, in both D1 and D2, the bank equivalent capacitance increases from a minimum value at stage $j = 1$ to a maximum value at stage $j = N$, with $N - 1$ discrete intermediate levels. For the same unit capacitance C_1 and number of stages N, D1 provides a

larger equivalent capacitance than D2 at each intermediate stage, while the two architectures converge to the same maximum capacitance at the final fully parallel stage. D1 therefore exposes higher capacitance values in the early and intermediate stages, whereas D2 starts from a much smaller equivalent capacitance and increases more steeply as the stage index grows, as shown in Fig. 5.14.

However, the equivalent capacitance alone does not determine the energy that can be extracted from a stage. During discharge, the energy delivered in each stage depends on both the stage capacitance and the voltage interval over which that capacitance is discharged. In D1, the transition-induced voltage spike becomes progressively smaller during discharge, so the additional energy provided by the post-switch voltage boost is strongest near the high-capacitance stages and weaker at lower stages. In D2, by contrast, every transition to a lower-capacitance stage doubles the bank voltage. Therefore, although D2 provides a smaller equivalent capacitance than D1 at the early and intermediate stages, each D2 discharge interval starts from a higher post-switch voltage, which partly compensates for the lower capacitance in terms of extractable energy. Since the extractable energy scales as $C(V^2_{\text{post switch}} - V^2_{SL})$, this increase in $V_{\text{post switch}}$ partly compensates for the lower capacitance of D2. As a result, D1 and D2 exhibit clearly different equivalent-capacitance distributions, as shown in Fig. 5.14, but more similar stage-energy availabilities, as shown in Fig. 5.13.

As a conclusive note, given a certain switching strategy, multiple capacitor dimensioning solutions can exist. We focus solely on one viable dimensioning solution per switching strategy, as the main contribution of this thesis lies in the autonomous bank reconfiguration, rather than the switching strategy or the capacitor bank dimensioning itself.

Load voltage range and reconfiguration spikes: once the bank is fully defined, it must be coupled to the desired load voltage range. In fact, the switching thresholds V_{SH} and V_{SL} must respect the load's voltage limits while also accounting for the voltage drops/spikes induced by reconfiguration. These post-switch voltage variations can be derived analytically from the capacitor-network topology, as summarized in Appendix A.1.

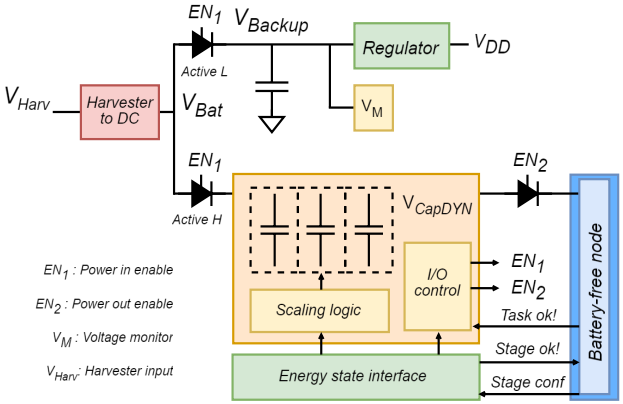
From this analytical framework, for the D1 bank the capacitance scale-up/down voltage drops/spikes are larger when transitioning between higher-capacitance stages. In particular, switching

to/from the $C_{\max}$ stage results in a halving/doubling of the main bank voltage, respectively. For the D2 bank, the halving/doubling of the bank voltage occurs at every transition.

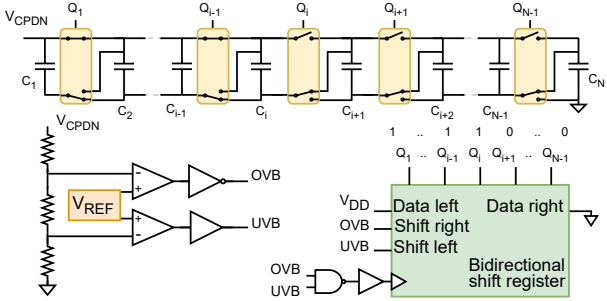
As a result, loads satisfying $V_{SH} \geq 2V_{SL}$ can fully exploit both capacitor bank designs, allowing charge accumulation up to the final stage. Conversely, loads with $V_{SH} < 2V_{SL}$ are not compatible with D2 and can only operate up to an intermediate stage $i < N$ in the D1 configuration, in order to avoid brownout. Most low-power integrated circuits fall in the former category, typically operating within the $[1.8, 3.6]$ V range. For loads in the latter case, *CapDYN* still achieves an equivalent minimum capacitance $C_{\min}$, at the cost of a reduced $C_{\max}$.

Overall, the analytical characterization of post-switch voltages enables CapDYN to be adapted to arbitrary load voltage ranges without requiring an additional output voltage smoothing or regulation stage.

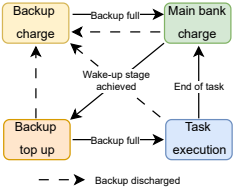
Interface to control signals: The last point of the main bank design is concerned with matching the switching logic to the switches' control signals needed to achieve the switching strategies presented in Fig. 4.1. The capacitor bank schematic and the control signal generation is shown in Fig. 4.2b: the connection of each capacitor is managed by two switches driven by the same control signal `Q_i`. We design the progressive scaling logic as an N-bit *bidirectional shift register* piloted by a window-type voltage monitor with two programmable thresholds coinciding with the capacitance scaling voltages, V_{SH}, V_{SL}. The voltage monitor provides two signals, `OVB` and `UVB`: `OVB` is asserted when the main bank voltage grows up to the upper scaling voltage, while `UVB` is asserted when the monitored voltage falls to the lower scaling threshold. As shown in Fig. 4.2b, `OVB` and `UVB` trigger the *right shift* of a high bit and *left shift* of a low bit of the bidirectional register bits, respectively, which will then control the main bank switches, responsible for the capacitance scale up/down. We refer to the N-bit bank control signal as `Q` or bitmask, alternatively. For both D1 and D2, the `Q` signals control the capacitor switches through comparators, which translate the digital voltage levels to the switch supply domain. By implementing different bitmasks, we can extend the progressive scaling logic to other designs.



(a) *CapDYN*'s high-level block diagram.



(b) D1 bank's schematic for the i-th stage. The window-type voltage monitor (left) triggers a left/right shift of the bidirectional shift register's bits. On the top, we can see the switch connections driven by the shift register's outputs, Q_i.



(c) FSM machine implemented by the I/O control section, which directs the input and output flow of both *CapDYN* and its backup capacitor.

Figure 4.2: *CapDYN* block diagram (a), bank schematics (b) and input/output logic (c). Adapted from [1].

4.2.3 Power supply design and input/output power flow management

CapDYN is designed around ultra-low power energy harvesters at the edge of remote energy scavenging, such as small photovoltaic panels, microbial fuel cells, thermoelectric, or even piezoelectric generators. Any node operating directly from such energy sources must rely on a *harvester-to-DC* conditioning circuit, as these sources produce voltage outputs that need to be either boosted, bucked or rectified to operational-level DC voltages. Accordingly, *CapDYN* is designed to receive the conditioned output of this source-specific front end rather than the raw harvester signal. It can therefore interface with any harvester whose output can be rectified, when required, and bucked or boosted to provide a DC voltage sufficient to charge the main bank to at least V_{SH}. This separates the design of the source-specific conversion electronics from that of the *CapDYN* architecture itself.

Power supply design: We power *CapDYN*'s load directly from the main bank's unregulated voltage, which is limited to the $[V_{SL}, V_{SH}]$ range, while we power *CapDYN*'s components by regulating the backup supply's voltage. Voltage regulation is deemed necessary to ensure the combinational logic's correct operation: the regulated voltage level is set to the minimum possible value to minimize power consumption. We dimension the backup capacitor to withstand the longest atomic operation of the application to which we are fitting the hardware, considering the worst-case scenario of no input power. Backup capacitor size must be kept as small as possible: as its size entirely depends on task execution time and *CapDYN*'s idle consumption, *CapDYN*'s components must be kept as low-power as possible, and time overhead due to backup capacitor charging must be considered when straying away from applications which do not entail short, intermittent tasks. A good rule of thumb is to keep the backup capacitor dimension below the capacitance of the first stage so that tasks with stage-1-level energy consumption are not excessively slowed down.

I/O control: the output of the selected harvester, which we consider to be already conditioned for the previously mentioned reasons, is the input of two separate but interdependent energy banks within *CapDYN* (Fig. 4.2a). One is the *main* supply (V_{CapDYN}), tapping into *CapDYN*'s main capacitor bank to power the intermittent node,

and the other is the *backup* supply (V_{Backup}) hinging on a backup capacitor to power *CapDYN*'s components. The power flow into and out of these two banks is controlled thanks to reverse-current blocking ideal diodes, piloted by the I/O control's two output signals, $\texttt{EN}_1$ and $\texttt{EN}_2$, as shown in Fig. 4.2a. In particular, the input power is multiplexed into either V_{Backup} or V_{CapDYN} by the use of two opposite-logic switches controlled by $\texttt{EN}_1$, to make sure only one switch is conducing at a time. On the other hand, the bank output is enabled by $\texttt{EN}_2$. The I/O control implements a finite state machine (FSM), shown in Fig. 4.2c. This FSM makes sure that power can flow into/out of the main bank solely when the backup capacitor can correctly supply *CapDYN*'s components and thus grant *CapDYN*'s self-scaling capabilities. The FSM also enables the main bank's output solely when *CapDYN* fully charges the capacitance level needed to execute the next task without power failures, which we term *wake-up* stage. By fully deactivating the load until useful work can be done, we avoid sleep mode power consumption on the load side and improve energy efficiency (Sec. 4.3.3). The FSM also safeguards *CapDYN*'s and the load's operation in the case of malfunctions: $\texttt{EN}_2$ interrupts the output in the case that either the backup or the main bank are too discharged to provide operative voltage to either the *CapDYN* infrastructure or the load. As will be explained in Sec. 4.4, this feature is useful during the testing phase and can be removed for marginal infrastructural power reduction once the development phase provides stable and predictable operation.

In its most basic functioning, the FSM ensures that the main bank begins charging when the backup is full. Then, when the wake-up stage is achieved, the backup is briefly topped up to ensure it can sustain *CapDYN* during execution. Once that happens, the load is activated and the task can execute. The load then power-gates itself by indicating to the I/O control when it has concluded its task, through signal `EOT`.

Operation under prolonged energy scarcity: The backup supply is designed not only to sustain the *CapDYN* infrastructure during task execution, but also to preserve its control state for as long as possible when harvested input power becomes temporarily unavailable. During such an outage, the load remains power-gated and the stored backup energy is progressively depleted by the quiescent

consumption of the control infrastructure.

If the available energy eventually becomes insufficient to sustain the digital control logic, the architecture must transition to a known safe configuration rather than leave the bank-control signals floating. In the proposed design, loss of the digital supply therefore forces the bank toward its minimum-capacitance configuration. The bank configuration remains preserved below the digital-logic operating voltage as long as the analog switches are still able to retain their channel state. If their supply also falls below its valid operating range, retention of the remaining bank energy can no longer be guaranteed.

Upon recovery of harvested power, the control circuitry is reinitialized through power-on reset to a user-defined default wake-up stage. The application state itself is retained in the load's nonvolatile memory, allowing the runtime to resume execution and reprogram the energy state required by the next task. These considerations introduce three additional component-selection requirements for prolonged-outage resilience: a low minimum operating voltage for the bank switches, a well-defined control state when the digital supply is absent, and a backup-supply capacity that balances state-retention time against cold-start latency. Their realization in the implemented prototype is described in Sec. 4.4.1.

Energy state interface: the key feature that enables energy-admissible task execution (Sec. 4.3.1) is CapDYN's *energy state interface*, which exposes the current energy state of the storage to the computing unit. The energy state interface allows the computing unit to request a specific energy level for the execution of the next task and to be notified when such level has been reached.

The scheduling contribution of this work lies in the use of energy-admissibility as a runtime safety invariant coupled with a programmable energy-storage interface. Unlike conventional threshold-based approaches, where the software reacts to a fixed voltage threshold exposed by the storage element, the proposed scheduler configures the storage state required by the next task before execution. The storage subsystem then autonomously accumulates energy and powers the computing unit only when the requested state has been reached. As a result, task admission is not based on continuous energy polling or prediction, but on a hardware-enforced correspondence between profiled task cost and discrete storage state.*wake-up register*, a *task match* section and five digital control signals exchanged between

CapDYN and the computing unit. Prior to task execution, the computing unit configures the wake-up register with a bitmask corresponding to the minimum capacitance stage required to execute the next task without power failure. The task match section then compares the requested wake-up stage with the current capacitance stage of the bank. When the requested stage is reached, the task match section asserts the *WKP* signal and enables the power supply to the computing unit through $\texttt{EN}_2$, allowing task execution to begin. Upon task completion, the load checks whether the successive task's energy is already available by piloting the next task register. If it is available, the computing unit immediately executes the successive task. If the energy isn't sufficient to continue executing, the computing unit asserts the *EOT* (End Of Task) signal, which tells to *CapDYN* to power-gate the load. This mechanism ensures that the computing unit is fully power-gated whenever insufficient energy is available, eliminating load-side sleep-mode consumption during these inter-task charging intervals and avoiding speculative execution.

The wake-up register is configured through a simple shift-based interface composed of initialization (*INIT*), clock (*CLK*), and read-back (*RB*) signals. During register reconfiguration, the *INIT* signal sets the wake-up register to a known, user-configurable default state, establishing a reference position for the active bit. Subsequent *CLK* pulses shift the active bit within the register to select the desired wake-up stage corresponding to the next task. The wake-up register is first initialized during *CapDYN*'s cold start through a power-on signal. This allows *CapDYN* to configure a default wake-up stage at which the load is first supplied with power, so that the runtime can initialize, as described in Sec. 4.3.3.

The *RB* (Read-Back) signal allows the computing unit to observe the state of the wake-up register and verify the position of the active bit. By sampling the *RB* line after each shift operation, the computing unit can confirm that the wake-up register has been correctly configured. This read-back mechanism increases the robustness of the hardware–software interface and guarantees correct alignment between task requirements and the configured energy state, without introducing additional monitoring overhead.

From the software viewpoint, the energy state interface exports two services: it allows the runtime to configure the minimum storage stage required by the next task, and it guarantees that the load is

powered only when that stage has been reached.

4.3 *CapDYN* software design

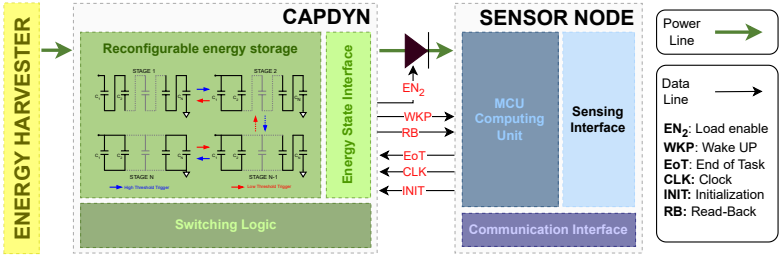


Figure 4.3: Overview of the proposed energy-aware sensor architecture. The system integrates an energy harvester, a dynamically reconfigurable energy storage subsystem (*CapDYN*), and a computing unit interfacing with generic sensing and communication modules. The energy state interface enables coordination between the storage subsystem and the computing unit to support energy-admissible task execution. Adapted from [2]. © 2025 IEEE.

Previous sections established the hardware foundations for energy-admissible execution. Specifically, *CapDYN* autonomously reconfigures its capacitance during charging and discharging, and exposes the current storage stage to the load through the energy state interface. What remains to be defined is the software layer that interprets this hardware state and translates it into a task-dispatch decision.

This section introduces that software layer. We present a runtime mechanism that associates each task with the minimum *CapDYN* stage required for its successful completion, and dispatches a task only when the energy state interface indicates that the corresponding stage has been reached. Unlike conventional intermittent runtimes, the proposed mechanism does not rely on online energy estimation, speculative execution, or rollback-based recovery. Rather, it uses the storage state exposed by *CapDYN* as a direct condition for task admission.

This design has two main implications. First, task execution is no longer speculative, since a task is released only after the energy required for its completion has been accumulated in the bank.

Second, power interruptions are effectively shifted outside task bodies and confined to inter-task charging intervals, thereby reducing runtime complexity and limiting state preservation requirements to minimal inter-task context. Figure 4.3 summarizes the resulting architecture, highlighting the interaction among the energy harvester, the reconfigurable storage subsystem, and the computing unit through the energy state interface.

The main contributions of this section are summarized as follows:

- We present the runtime counterpart of *CapDYN*, namely an energy-aware scheduling framework for intermittent batteryless systems equipped with dynamically reconfigurable storage.
- We introduce an energy-admissibility mechanism based on *CapDYN*'s energy state interface, whereby task release is conditioned on the storage stage required to guarantee task completion.
- We show how offline task-to-stage mapping enables the systematic exploitation of CapDYN across heterogeneous workloads, avoiding oversized storage provisioning and preventing mid-task failures due to insufficient stored energy.

4.3.1 Execution Model and Definitions

We consider a generic batteryless sensing system composed of an energy harvesting source, an energy storage element, a microcontroller, and a set of peripheral devices. Application behavior is expressed as a finite sequence of tasks, where each task represents an atomic sensing, processing, or communication operation. Tasks are assumed to be non-preemptive and are executed sequentially.

Let $E_s(t)$ denote the energy stored in the system at time t, and let T_i be a task associated with a profiled energy cost E_i. Energy profiling is performed offline by repeatedly executing each task under controlled conditions and measuring its average energy consumption, as explained in Sec. 4.4.4. Based on this information, we define the following.

Definition 1 (Energy-Admissible Task). A task T_i is said to be *energy-admissible* at time t if:

$$E_s(t) \geq E_i + E_{margin}, \tag{4.4}$$

where E_{margin} accounts for system-level overheads, energy conversion losses, and residual variability in task execution. Under this condition, the energy stored in the system is sufficient to sustain the entire execution of task T_i without interruption. Tasks that are not energy-admissible are not dispatched and remain pending until sufficient energy is accumulated.

The proposed scheduler enforces energy-admissibility as a strict runtime invariant coupled with a programmable energy-storage interface. Before dispatching a task, the scheduler admits it for execution only if the admissibility condition is satisfied. Unlike conventional threshold-based approaches, where the software reacts to a fixed voltage threshold exposed by the storage element, the proposed scheduler configures the storage state required by the next task before execution. The storage subsystem then autonomously accumulates energy and powers the computing unit only when the requested state has been reached. As a result, task admission is not based on continuous energy polling or prediction, but on a hardware-enforced correspondence between profiled task cost and discrete storage state. This approach ensures *zero-failure execution with respect to energy availability*: once a task is admitted, its completion is guaranteed in the absence of hardware faults. Partial executions and repeated restarts due to energy depletion are avoided, thereby improving the fraction of harvested energy that contributes to successful task execution and increasing effective energy utilization. The execution model is *deterministic* with respect to energy availability: given the same sequence of harvested energy events and task energy profiles, task execution order and completion behavior are fully determined by the scheduler, simplifying system analysis and performance evaluation.

The proposed execution model assumes the presence of an energy storage subsystem capable of exposing a fine-grained representation of its energy state to the scheduler. In the implemented system, this interaction is realized through a hardware–software co-design based on *CapDYN*. The scheduler configures the minimum storage state required for the next task based on its profiled energy cost. Task execution is enabled only when the storage subsystem signals that the required energy level has been reached. Upon task completion, the scheduler explicitly terminates execution and CapDYN power-gates the load, preventing further energy consumption until the next task becomes admissible. Importantly, the scheduler does not rely

on continuous energy monitoring or prediction. Energy awareness is instead achieved through discrete energy states exposed by the storage subsystem, allowing the scheduler to operate with negligible runtime overhead.

It is worth noting that energy-admissibility does not replace higher-level scheduling policies. Rather, it acts as an admission-control layer that determines whether a selected task can be safely executed with the currently available stored energy. The choice of the next task can still be performed by the application logic or by a higher-level policy, such as a priority-based, latency-aware, or mixed-criticality scheduler. Once a task is selected, the energy-admissibility condition only determines whether sufficient energy is available to dispatch it. In this sense, energy-admissibility constrains execution based on energy feasibility, but does not prescribe the ordering of the task queue.

4.3.2 Assumptions, Scope, and Limitations

The proposed execution model relies on several assumptions. First, task energy costs are assumed to be known a priori through offline profiling. This does not imply that the instantaneous power consumption of a task is perfectly deterministic. Rather, each task is associated with a conservative energy bound obtained from repeated measurements under representative operating conditions, as described in Sec. 4.4.4. The admissibility margin in Eq. (4.4) is introduced to account for residual variability due to sensing conditions, peripheral behavior, communication overhead, temperature, and measurement uncertainty. This conservative design may lead to storage overprovisioning, but it prevents task failures due to moderate runtime variations.

Second, tasks are assumed to be *atomic* and *non-preemptive*. This assumption is consistent with the class of battery-free sensing workloads considered in this work, where many operations involve external peripherals, such as sensors, cameras, or radio transceivers. Interrupting these operations may leave the peripheral in an inconsistent state, invalidate a transaction, or require additional recovery logic. For these tasks, preemption would introduce extra latency and energy overhead without providing a clear benefit. Purely computational workloads could instead be decomposed into smaller atomic tasks, while checkpointed or preemptive extensions could be sup-

ported by treating execution segments as energy-profiled units, provided that the energy required to save and restore the execution state is explicitly included in the admissibility condition.

Temporal constraints and mixed-criticality objectives can be supported by exploiting the programmable energy-state interface to reserve stored energy for time-critical operations. In this case, the task-selection policy can defer non-critical tasks until a reserved charge level is reached, or prevent lower-priority tasks from consuming energy states allocated to urgent sensing, communication, or safety-relevant operations. Different task classes can also be assigned different admissibility margins or required energy states, so that critical tasks are admitted only when a more conservative energy reserve is available. However, such reservations remain bounded by the energy actually harvested by the system. In battery-free intermittent systems, harvested energy is still the primary execution constraint: if sufficient energy is not accumulated before a deadline, no scheduling or reservation policy can guarantee task completion. Therefore, the proposed model can support priority-aware and mixed-criticality execution by reserving charge and controlling task admission, but it does not provide hard real-time guarantees under arbitrary harvesting conditions. The evaluation in this work focuses on energy feasibility and reliability, while latency optimization, fairness-aware scheduling, and explicit energy-reservation policies for time-critical workloads are left to future work.

Finally, the model assumes that the energy storage subsystem exposes a stable and monotonic representation of the available stored energy. In the proposed implementation, this information is provided by the discrete stages of *CapDYN*. The energy associated with each *CapDYN* stage is estimated from the nominal capacitance values. Therefore, conservative margins are required in the stage-to-energy mapping, because capacitor tolerances, temperature variations, and aging can affect the actual energy available at each stage. This limitation can be mitigated through experimental calibration of the capacitor bank and by adapting the LTspice model to measured component parameters, so that the simulated stage energies more closely reflect the behavior of the implemented hardware under representative operating conditions. Significant changes in the operating environment may still require re-profiling of task costs or recalibration of the energy-state mapping.

In the present implementation, this re-profiling remains an offline

procedure, and the number of physically available *CapDYN* stages is fixed at design time; possible online adaptation to initially unknown or changing workloads is discussed in Sec. 5.3.

Although this work implements the proposed model using *CapDYN*, the scheduling principle is not specific to the internal circuitry of *CapDYN* itself. The required hardware support is an energy storage subsystem capable of exposing discrete energy states and enforcing load gating. Therefore, the proposed model should be interpreted as a hardware–software co-design abstraction for systems with explicit energy-state feedback, rather than as a fully hardware-agnostic scheduling framework.

4.3.3 *CapDYN* Task-based, Energy-Aware Runtime

CapDYN utilizes an energy-aware, task-based intermittent runtime model. The *CapDYN* runtime executes each task by considering if the task is *energy admissible*, to avoid power failures. To this end, *CapDYN* requires an offline task-energy profiling process, described in Sec. 4.4.4, to derive a *task-to-stage* mapping that associates each task with the minimum *CapDYN* stage required for its successful completion.

Using this mapping, the runtime programs the energy state interface with the stage required by the next task. If the storage has already reached that stage, the task is executed immediately. Otherwise, the runtime terminates its current activity and allows the hardware to power-gate the load until the requested stage is reached. In this way, the runtime delegates energy-state tracking to *CapDYN* and only resumes execution when the hardware has established that the next task is energy-admissible.

Figure 4.4 summarizes the runtime flow. When the load is first activated, the runtime initializes the first task by programming the stage associated with its task-to-stage mapping. After each task completes, the runtime configures the stage required by the subsequent task and determines, through the energy state interface, whether execution may continue immediately or must be suspended until further energy has been accumulated. This process allows the runtime to progress through the task sequence without incurring mid-task power failures.

A key consequence of this design is that the load does not need

to perform direct state-of-charge monitoring. The only energy-awareness required by the runtime is whether the storage stage needed by the next task has been reached, and this information is already encoded in the hardware interface. As a result, *CapDYN* avoids both periodic polling and interrupt-driven wake-up mechanisms for energy-state tracking, thus solving the reactivity-efficiency trade-off.

Polling-based approaches incur a twofold limitation. First, they consume additional energy by repeatedly activating the load to sample the storage state. Second, they require the designer to select a monitoring period, which introduces an inherent trade-off between reactivity and efficiency: frequent polling improves responsiveness but increases overhead, whereas sparse polling reduces overhead at the cost of delayed reaction to changes in the harvested power. As a result, the system is neither continuously reactive nor continuously reconfigurable, and may fail to exploit short-lived energy bursts [22, 23].

Interrupt-driven approaches improve upon polling by avoiding periodic state checks, but they still require the load to remain responsive to state-of-charge notifications [91]. This typically implies keeping the computing unit in a low-power sleep state rather than fully power-gating it, thereby incurring residual energy consumption while waiting for sufficient energy to accumulate.

CapDYN removes both limitations. Since energy-state detection and bank reconfiguration are handled autonomously in hardware, the load does not need to periodically inspect the storage state nor remain awake to service energy-driven interrupts. Instead, it is fully power-gated until the required energy state is reached, while the storage remains instantly reactive to variations in the harvested input. This allows *CapDYN* to combine immediate responsiveness to incoming energy bursts with negligible control overhead.

4.4 Hardware and Software Implementation

In this section, we describe the hardware implementation of *CapDYN*, its lightweight runtime, and the LTspice model used for subsequent evaluation. We implement both D1 and D2 capacitor-bank topologies, together with the control, switching, power-flow, and

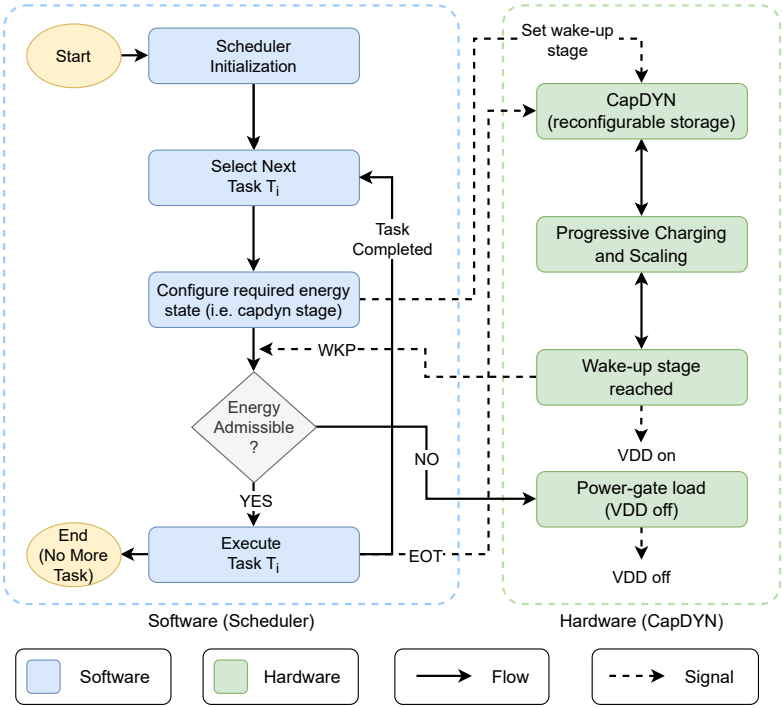


Figure 4.4: Energy-admissible task execution flow implemented through a hardware–software event-driven interaction between the scheduler (left) and CapDYN (right). The scheduler configures the required energy state for the next task and initiates execution only when CapDYN signals that the requested state has been reached. Task execution and power gating are coordinated via event signals. Reproduced from [2]. © 2025 IEEE.

monitoring blocks introduced in Sec. 4.2. Given the load-voltage range and the profiled task-energy requirements, the capacitor bank can be dimensioned analytically using the stage-capacitance and post-switch-voltage models of Sec. 4.2.2 and Appendix A.1. The runtime exploits the energy-state interface by dispatching tasks only when the corresponding energy state is available. Finally, we present the LTspice model of *CapDYN*, which complements the analytical dimensioning and enables the later application-level comparison of D1 and D2 topologies with different total numbers of stages under realistic harvested-power traces. The hardware presented here is characterized experimentally in Sec. 5.1.4, while the LTspice-based application-level comparison is reported in Sec. 5.2. Both the hardware design files and the LTspice simulator are released as open-source [93].

4.4.1 Hardware Implementation

In designing *CapDYN*'s schematic, we sought the lowest power-consuming solutions to minimize *CapDYN*'s energy overhead and extend its autonomy upon power failure. The switches chosen for the main capacitor bank are Analog Devices' ADG801 (SPST) and ADG819 (SPDT). Their 1 nA current consumption ensures minimum burden on the backup supply, while their low on-resistance (0.4 and 0.8 Ω max, for the ADG801 and ADG819, respectively) minimizes the charging/discharging energy loss and capacitor bank imbalances. These capacitor bank switches conduct equally well in both directions, allowing both charging and discharging operations. Unlike the capacitor bank switches, the power input/output's ideal diodes (Fig. 4.2a) need to provide reverse current protection: we thus use Vishay's reverse current blocking load switches SIP32431, which consume only 10 pA and introduce 130 nA of leakage when reverse biased.

Two Analog Devices' LTC2934-2 voltage monitors are then used for piloting the progressive scaling logic and the I/O power flow management described in sec. 4.2.3. Each voltage monitor has a quiescent current of 500 nA and features a resistive divider at its input, through which we set two threshold levels on the tracked voltage. The selected voltage monitor features ultra-low leakage at the two control inputs (2 nA combined). The maximum resistor divider size is set to achieve a 1% voltage detection error due to

the leakage of the control inputs. The divider is dimensioned so that the divider current is 100 times larger (so, 0.2 µA) than the inputs' leakage at the lowest monitored threshold. An alternative approach is to implement the window-type voltage monitor directly from discrete components: by using Texas Instruments' nano-power TLV3691 comparators and a REF35 0.4 V, 650 nA fixed voltage reference, we can achieve a comparable quiescent power draw with respect to the integrated solution.

The first voltage monitor is configured to track the capacitor bank voltage so that it matches the target load range. We evaluated the system assuming the load voltage range is the one typically found in ultra-low power computing units (1.8-3.6 V). We thus set the two thresholds at $V_{SL} = 1.8$ and $V_{SH} = 3.65$ V, leaving a small 50 mV margin to ensure correct scaling even in the case of the largest voltage spikes/drops (Sec. 4.2.2). Whenever the bank voltage grows up to 3.65 V (or falls to 1.8 V), the voltage monitor triggers capacitance scale-up (or down) thanks to the progressive scaling circuit.

The second voltage monitor is set to track the backup capacitor voltage: as the backup capacitor powers all components making up the *CapDYN* infrastructure, maximum and minimum backup capacitor voltage values must be enforced so that power input/output flow into/out of the bank is allowed only when the backup can correctly power the *CapDYN* infrastructure (Sec. 4.2.3). In fact, *CapDYN*'s digital components' supply line is obtained by stepping down the backup capacitor voltage to 3 V, through Texas Instruments' TPS62840DLCR 60-nA quiescent current buck converter. The voltage output value is configurable in 100 mV steps through a programming resistor: the chosen 3 V value represents the minimum voltage supply of the selected digital components (amongst which, the Texas Instruments' CD4000 family has the highest minimum voltage). Analog components such as the bank's switches are directly powered from the unregulated backup capacitor voltage. In order to further minimize the analog switches' power consumption, the `Q` signals generated from the progressive scaling mechanism, were voltage shifted to the voltage supply of the switches (the backup capacitor), through the use of Texas Instrument's nanopower TLV3691 comparators.

The lower backup voltage threshold is selected to prevent the clamping diodes of the analog switches from entering conduction

during normal operation. The switches' supply, and so the backup voltage, should not drop below the switches' maximum input voltage (V_{SH}) plus their protection diodes' forward drop (0.3 V), which in this implementation amounts to 3.85 V. The buck converter successfully provides a stable 3 V output with a backup capacitor input voltage above 3.1 V, so the chosen lower backup voltage threshold also ensures correct digital components powering. On the other hand, the upper backup voltage threshold is set to 5.5 V, which is the maximum voltage value accepted by the chosen ADG switches.

The lower backup voltage threshold therefore defines the intended operating limit of the backup supply. During a prolonged absence of harvested input power, however, the backup capacitor continues discharging after this limit has been crossed. Because the backup capacitor is typically dimensioned to be smaller than the main bank and directly supplies the quiescent consumption of the control circuitry, its voltage is expected to decrease faster than the main-bank voltage. If the backup voltage approaches the main-bank voltage, the input-protection diodes of the bank analog switches can become forward biased. Since the switches are powered directly from the backup capacitor, this occurs when a switch terminal rises approximately 0.3 V above the corresponding supply rail. A current path is then established between the main bank and the backup supply, causing the two energy stores to discharge effectively in parallel.

Although this coupling is undesired during normal operation, it extends the tolerance of the implemented architecture to prolonged outages in two ways. First, while the backup voltage remains high enough for the TPS62840 to sustain the 3 V digital rail, energy remaining in the main bank can contribute to supplying the control infrastructure. The parallel discharge therefore extends the interval over which the digital control state can be retained. Second, the same coupling supports the supply of the bank analog switches, allowing their channel configuration to be maintained for longer and thereby delaying the loss of the energy remaining in the main bank. The 3.1 V and 1.8 V limits are determined by the regulator and analog switches, respectively, and apply to the backup capacitor voltage independently of whether coupling between the two storage elements occurs; the coupling only delays the time at which these limits are reached.

If the outage persists until the backup voltage can no longer sus-

tain the regulated digital supply, the digital control logic is powered off and its volatile state is lost. The bank-switch control signals are pulled low rather than being left floating, forcing the main bank into its minimum-capacitance configuration. Although the digital control logic is no longer active, the analog switches can continue to maintain this configuration, and therefore preserve the energy remaining in the bank, as long as their supply remains above their minimum operating voltage of approximately 1.8 V.

If the outage continues until the backup supply, and the main bank when the two have become coupled, fall below approximately 1.8 V, the channel state of the analog switches can no longer be guaranteed. The capacitor-bank configuration can therefore no longer be reliably maintained, and the remaining stored energy may be lost. When harvested power subsequently becomes available, the power-on-reset mechanism initializes the wake-up register to its user-configurable default stage, normally the first stage. Once the load is activated, the runtime recovers its task-control state from non-volatile memory and programs the wake-up stage required by the next task, allowing application execution to resume.

The duration of this retention interval depends on the backup capacitance. In the present implementation, the backup capacitor is dimensioned to sustain the longest atomic task and is kept small to limit cold-start latency. Increasing its capacitance would improve tolerance to prolonged outages at the cost of a longer cold start. Future implementations could instead improve state retention by adopting bank-switch designs that preserve their channel state without an actively powered control supply, for example through state-retaining techniques such as the latch-capacitor-controlled switches employed by Capybara [21], provided that the required leakage, on-resistance, and operating-voltage constraints can simultaneously be satisfied.

The I/O flow control relies on the two voltage monitor signals to follow the FSM defined by Fig. 4.2c. Ultimately, for the progressive scaling logic and the energy state interface and the I/O control, both showcased in Fig. 4.2a, we relied on Texas Instruments' ultra-low power digital components: CD4000 logic gates (CD4011 NAND, CD4071 OR, CD4081 AND) and flip flops (CD4013).

Within the I/O control, load powering can be interrupted by either the bank discharging below 1st stage, V_{SL} or by the backup discharging below its lower voltage limit. This feature is useful

in testing and dimensioning phases and shouldn't be necessary in normal operation, where load powergating is issued only by the load itself through the `EOT` signal. Users can freely choose to remove this feature for a marginal infrastructural power reduction, as the two additional interruption signals are ORed with the `EOT` signals.

The complete schematics for the scalable 3-, 4-, and 5-stage implementations of both the D1 and D2 banks are released as open-source in the accompanying repository [93]. In particular, the repository includes the *CapDYN* control schematic, the D1 bank schematic, the D2 bank schematic and the corresponding PCB layouts. Fig. 5.2 and Fig. 5.3 show the PCB layouts implementing the *CapDYN* control logic with the D1 bank and the D2 bank, respectively, while Fig. 5.4 shows the assembled prototype used during the hardware characterization.

4.4.2 LTspice simulator

We fully simulate *CapDYN* in LTspice, with the intention of providing a hardware-software co-design tool that minimizes the number of prototype iterations to reach the best fit between a certain task set and *CapDYN*. With this simulator, we can easily include both the profiled average values and execution times for the task set we aim to supply with *CapDYN*, as well as inject real-world power input traces coming from a target energy harvesting source. This allows the user to fine-tune the *CapDYN* parameters for the task at hand, and test overall system behaviour under particular power input conditions. For maximum simulation accuracy, we use all the available supplier-provided LTspice models of the selected components. We favoured transistor-level models to get an accurate estimation of the component's power consumption, only considering behavioural models if the transistor-level models were not available. To integrate the behavioural models, we sourced their datasheet power consumption from the relevant supply line. Only in the case of the Vishay load switches, the SIP32431, no model was available: we developed a model that takes into account their reverse leakage and channel resistance effects. The LTspice model extends the hardware measurements to application-level design-space exploration. The complete system-level evaluation, including the target battery-free application, input power traces, task-to-stage mapping, baseline comparisons, and application-level performance metrics, is

```
#include "capdyn.h"
//Task declaration
Task task_t1() {
    ...
    transition_to(task_t2);
}
Task task_t2() {
    ...
    transition_to(task_t3);
}
Task task_t3() {
    ...
    transition_to(task_t4);
}
Task task_t4() {
    ...
    transition_to(task_t1);
}
...
//Defining Required Tasks' Stages
set_stage(task_t1, Stage2);
set_stage(task_t2, Stage3);
set_stage(task_t3, Stage1);
set_stage(task_t4, Stage4);
...
Entry_Task(task_t1);
```

Figure 4.5: Example application code using the *CapDYN* runtime. Adapted from [1].

presented in Sec. 5.2.

4.4.3 Software Implementation

We implemented the *CapDYN* runtime as a lightweight C library supporting task-based applications [107, 109, 108] and interfacing with the *CapDYN* hardware through the finite-state machine shown in Fig. 4.4. The microcontroller powered by *CapDYN* executes this runtime, which is responsible for task scheduling, enforcement of energy admissibility, and coordination with the energy state interface. Accordingly, a task is dispatched only when the current storage state satisfies the admissibility condition defined in Sec. 4.3.1. Our target platform is an MSP430FR [117] series microcontroller, which is widely used in intermittent systems and features ultra-low power operation and integrated non-volatile FRAM. The *CapDYN* run-

```
1  #include "capdyn.h"
2  ...
3  int main(){
4      ...
5      if (!system_initialized){
6          curr_task = initial_task;
7          system_initialized = SET;
8      }
9      transition_to(curr_task);
10 }
11 void transition_to(task_t *next_task)
12 {
13     uint8_t required_stage = stage_map[next_task->idx];
14     next_task_setup(required_stage);
15     int wake_up_stage = gpio_get_value(WKP_PIN);
16     if (wake_up_stage == 1){
17         exec_task(next_task);
18     }
19     else if (wake_up_stage == 0){
20         curr_task = next_task;
21         gpio_toggle(EOT_PIN);
22     }
23     ...
24 }
```

Figure 4.6: Code snippet that illustrates how system initialization and task transitions are managed within the *CapDYN* runtime. The `transition_to`() function sets up the HW register for the next task (line 14), checks if the required stage is achieved (line 15) and either executes the next task (line 16-20) or orders *CapDYN* to cut the load's supply (line 21-23). Adapted from [1].

time can be adapted to work on any platform that has non-volatile memory such as FRAM or MRAM.

In this subsection, we discuss the implementation of flagship functions of the software library, for an N-stage *CapDYN*. Fig. 4.5 presents an example application code that includes four tasks. The *CapDYN* runtime uses three main C macros to build an application: `Task`, `set_stage`, and `Entry_Task`. The `Task` macro is used to define tasks. Each task is a void function and has a corresponding data structure of type `task_t`, which holds non-volatile task control variables, such as the task's address and ID. The transition to the next task is managed by the `transition_to` macro at the end of each task. The `set_stage` macro sets the task's required stage

within the `stage_map`. Finally, the `Entry_Task` macro declares the entry task (`initial_task` in Fig. 4.6) during system initialization.

Fig. 4.6 presents the code snippet of the main entrance and the function responsible for the task-to-task transition (`transition_to`), which is invoked upon each wake up and after each task execution, to implement the set/check/execute loop of Fig. 4.4. When the *CapDYN* hardware triggers the load for the first time with the default wake-up stage, which is user configurable, the *CapDYN* runtime initializes the first task in the lines 5 to 8 and calls the `transition_to` function. The `transition_to` function is responsible for checking the stage requirement of the task and the current stage of the *CapDYN* hardware: it either executes the task if the current stage is sufficient, or it cuts the node's power if the current stage is not sufficient. In more detail, this function obtains the required stage of the task utilizing the `stage_map` array and the task ID at line 13. Then it writes the required stage to the wake-up register in line 14. *CapDYN*'s `WKP_PIN` informs the runtime if the required wake-up stage has been achieved in the capacitor bank: thus, in line 15, `transition_to` checks the `WKP_PIN`. If the required stage has been achieved, the runtime executes the task in line 17: at the end of the task's execution, `transition_to` is newly called to check if the successive task can in run immediately or not. If the required stage has not been achieved, *CapDYN* saves the task information to the `curr_task` in line 20 and sends an `EOT` - end of task - signal to *CapDYN* hardware to cut the power of the load in line 21. Thus *CapDYN* prevents wasting power by turning off the MCU when waiting for the charging of the required stage. When the required stage is achieved, the *CapDYN* hardware newly powers the load and invokes `transition_to` to check again if the saved task in `curr_task` in line 20 can execute. Thus, *CapDYN* executes all tasks considering their energy requirements and preventing power failures during their execution.

4.4.4 Energy Profiling of Tasks and Task-to-Stage mapping

As explained in Sec. 4.1.1, each stage represents a certain state of charge of the bank. Each stage is also connected to the tasks it can sustain without power failure. This task-to-stage mapping is obtained by first profiling the power consumption of the task set,

and then designing a *CapDYN* instance to best match each task's measured energy consumption. Offline profiling is a fundamental step, as we design *CapDYN* to perfectly sustain a certain task set. The more a stage matches the energy demands of a certain task, the more *CapDYN* will reduce charging times and improve reactivity. The profiling must be executed by considering both the run-to-run energy consumption variability and the variable voltage that *CapDYN* supplies to the load, in order to avoid power failures even in worst-case scenarios. To account for variable voltage during profiling, we fix the load voltage at the load's upper voltage threshold - which is where the loads consumes more power. The load's upper voltage threshold is set by *CapDYN*'s upper scaling threshold, V_{SH} (Sec. 4.2.1). Then, to account for run-to-run variability, we collect multiple (e.g. enough to have a statistical sample) energy consumption measurements, ultimately considering the largest consumption linked to a successful task execution as reference value for the dimensioning of *CapDYN*. This will slightly overprovision *CapDYN* with respect to the average run, but it will ensure faultless execution for higher-consuming runs. On the other hand, designers can analyze the energy consumption of each task by using software tools such as the Energy Trace Tool of Code Composer Studio. Additionally, they can employ energy profiling tools like Echo [118] for more accurate energy measurements. Once the energy consumption profile of each task is obtained, it can be matched to the energy level of the *CapDYN* capacitor bank's stages.

Once the energy profiling step is concluded, a suitable *CapDYN* dimensioning must be derived. The relevant parameters are the bank topology (here, D1 or D2), the unit capacitance C_{unit}, and the number of stages N. By treating C_{unit} as a free design parameter, the minimum energy available at each stage can be expressed analytically as a function of C_{unit} for a given *CapDYN* instance. More formally, if $\mathcal{S} = \{1, \ldots, N\}$ denotes the set of stages, we define the stage-energy function

$$E_s^{\min}(C_{unit}) \qquad \forall s \in \mathcal{S}, \tag{4.5}$$

where $E_s^{\min}$ is derived from the post-switch voltage reached after each stage transition, whose computation is summarized in Appendix A.1.

Let $\mathcal{T} = \{\tau_1, \ldots, \tau_M\}$ be the application task set, and let

$$\mathbf{E} = \{E(\tau_1), \ldots, E(\tau_M)\} \tag{4.6}$$

be the set of profiled task energies. We choose C_{unit} such that the largest stage can sustain the most energy-demanding task, i.e.

$$E_N^{\min}(C_{unit}) \geq \max_{\tau \in \mathcal{T}} E(\tau). \tag{4.7}$$

Subsequently, each task $\tau \in \mathcal{T}$ is associated with the minimum stage able to sustain it, yielding the *task-to-stage* mapping

$$\phi : \mathcal{T} \to \mathcal{S}, \qquad \phi(\tau) = \min\{s \in \mathcal{S} \mid E_s^{\min}(C_{unit}) \geq E(\tau)\}. \tag{4.8}$$

Finally, we evaluate how well the resulting *CapDYN* dimensioning fits the task set through an *overprovisioning* metric, which quantifies the excess stage energy assigned to each task with respect to its profiled energy demand.

$$\mathrm{OP} = 100 \cdot \frac{\sum_{\tau \in \mathcal{T}} \left(E_{\phi(\tau)}^{\min}(C_{unit}) - E(\tau) \right)}{\sum_{\tau \in \mathcal{T}} E(\tau)}, \tag{4.9}$$

which measures the total excess energy provisioned by the selected stage assignment relative to the total requested task energy.

A more refined dimensioning strategy could formulate the choice of (T, C_{unit}, N) as an optimization problem aimed at minimizing the overprovisioning metric, rather than selecting C_{unit} solely so that the largest stage sustains the most energy-demanding task. We leave the investigation of such optimization-based design-space exploration to future work.

4.5 Summary

This chapter presented the design and implementation of CapDYN, including the progressive capacitance-scaling mechanism, capacitor-bank architectures, energy-state interface, task-based runtime, and task-to-stage mapping procedure. The resulting architecture combines autonomous storage reconfiguration with runtime-visible energy states. The following chapter evaluates the implemented hardware and quantifies the resulting system-level benefits under representative battery-free workloads and harvested-power traces.

Chapter 5

Experimental and System-Level Evaluation of CapDYN

Parts of this chapter are adapted from the author's previously published work on CapDYN [1] and from the related IEEE SENSORS work on zero-failure task execution in batteryless sensors [2].

This chapter evaluates the *CapDYN* architecture developed in Chapter 4 at both circuit and application level. The evaluation connects the physical behavior of the reconfigurable storage with its impact on intermittent execution.

We first characterize the hardware prototype to validate the autonomous stage-switching mechanism, quantify the losses introduced by the capacitor-bank topologies, and measure the infrastructural power required by the control circuitry. Hardware measurements on D1 and D2 prototypes confirm autonomous stage progression during charging and discharging, stable load-voltage operation across reconfiguration events, and a strong reduction of leftover energy. Compared with a fixed-capacitor baseline, which leaves 23.67% of the input energy unavailable below the load operating range, the measured *CapDYN* configurations reduce leftover energy to 0.16%–5.57%, while keeping the control infrastructure below 10 μW.

These measurements establish the physical tradeoffs of *Cap-*

DYN. Dynamic reconfiguration reduces leftover energy and exposes intermediate energy states, but it also introduces infrastructural and bank-side parasitic losses. As a result, the total system-level efficiency of the current *CapDYN* implementation can be lower than that of a single-capacitor benchmark when the storage is evaluated only as a full charge–discharge buffer, mainly because of infrastructural consumption.

However, this efficiency-only characterization does not capture the main advantage of *CapDYN*. The greatest impact of the proposed system emerges when the storage is coupled to a load that exploits the *CapDYN* energy-state interface to perform failure-free, energy-admissible execution. In this operating mode, the load does not need to wait for the capacitor bank to reach its maximum-energy state; instead, each task can execute as soon as the corresponding stage provides sufficient energy. Therefore, even when the measured system-level efficiency is penalized by infrastructural overhead, *CapDYN* can still improve useful task execution by avoiding failed attempts and by making intermediate stored-energy states available to the scheduler.

We then use a complete LTspice model to evaluate how these hardware properties affect a realistic battery-free sensing workload. The LTspice evaluation complements the hardware measurements in two ways. First, it couples the validated *CapDYN* behavior to profiled task workloads, an energy-admissible scheduler, and realistic harvested-energy traces, showing how the energy-state interface translates the circuit-level behavior into application-level performance. Second, it extends the design-space exploration to 6- and 7-stage configurations, which are costly and bulky to realize with the present discrete-component implementation but are useful for studying the tradeoff between finer energy granularity and increased infrastructural power consumption. The evaluation also considers two representative task-energy distributions: an exponential task set, which better matches the approximately exponential energy distribution exposed by the staged banks, and a linear task set, which provides a less favorable matching case.

The application-level evaluation compares different *CapDYN* bank topologies and stage counts with fixed-capacitor and ablated reference systems. Against an energy-unaware fixed-capacitor system, which may start tasks without sufficient stored energy, *CapDYN* improves effective efficiency by 27.09–34.10% for the exponential task

set and by 13.30–19.56% for the linear task set, while increasing successful-task throughput by up to 41.67%. Even against energy-aware fixed-capacitor baselines, which already avoid failed executions but expose only a single storage threshold, *CapDYN* improves effective efficiency by up to 11.54% and advances first successful task completions by up to 41.03%. The ablation study further shows that these gains arise from the interaction between the reconfigurable bank and the energy-aware runtime: the bank exposes intermediate energy states, while the runtime turns them into earlier energy-admissible task executions.

The results of both the hardware and LTspice analyses also highlight a clear path for improvement. At bank level, *CapDYN* already achieves higher energy efficiency than the single-capacitor benchmark, primarily because dynamic reconfiguration reduces the amount of energy left unusable below the load voltage threshold. The remaining gap at system level is therefore mainly due to the infrastructural consumption required by the present discrete-component implementation. Reducing this overhead would allow *CapDYN* to more fully translate its bank-level efficiency advantage into system-level gains, while further amplifying its throughput and responsiveness benefits in battery-free applications.

The chapter is organized as follows. Sec. 5.1 presents the hardware evaluation of the implemented D1 and D2 capacitor-bank prototypes, validating autonomous reconfiguration and quantifying infrastructural consumption, bank-level losses, leftover energy, and overall efficiency. Sec. 5.2 extends the evaluation to the application level by coupling the LTspice model to profiled workloads, realistic harvested-power traces, and the energy-admissible runtime. The analysis compares different bank topologies and stage counts and quantifies their impact on effective efficiency, successful task execution, and responsiveness. Finally, Sec. 5.3 discusses the limitations of the current implementation and future research directions, while Sec. 5.4 summarizes the main findings of the chapter.

5.1 Hardware evaluation

The hardware evaluation focuses on the physical behavior of the *CapDYN* storage subsystem. Its goal is not yet to evaluate task-level performance, but rather to verify that the capacitor bank re-

configures as expected, that the staged energy levels are available to the load, and that the main loss mechanisms can be quantified experimentally. These results provide the basis for the subsequent LTspice evaluation, where the same storage behavior is coupled to the application scheduler and exercised under realistic harvested-energy traces.

5.1.1 Efficiency measurement setup

The efficiency measurement test bench can be seen in Fig. 5.1. The purpose of this test bench is to measure the charge-discharge efficiency as a function of input/output current of a 3,4,5 *CapDYN* D1 and D2 implementation and compare it to a single capacitor benchmark.

Charge and discharge in *CapDYN* are performed using two Source Measure Units (SMUs), both configured as current sources set to the same current level. The source SMU (Keithley 2470) simulates the output of the harvester-to-DC stage, while the sink SMU (Keithley 2450) represents the load, which is activated only when the requested stage is reached. In this case, the wake-up stage is set to the last stage, so that the bank achieves full charge and discharge. The SMUs are triggered through their digital I/O interface by the $\texttt{EN}_2$ signal produced by *CapDYN*. An external ADG819 switch powered by the digital I/O's power line is used to reference $\texttt{EN}_2$ to the digital I/O's logic levels.

CapDYN toggles $\texttt{EN}_2$ when the requested stage is charged: at that point, the source SMU stops sourcing power, while the sink SMU starts sinking it. As no load is present to prevent full *CapDYN* depletion through the `EOT` signal, the sink SMU stops drawing current when *CapDYN* toggles $\texttt{EN}_2$ upon full discharge (V_{SL}, 1st stage). Each operating point was repeated five times to assess measurement repeatability. After each run both the bank and the backup capacitor were fully discharged through a 1 kΩ resistor.

Voltage measurements are collected at 1.000 kHz through a LabVIEW program controlling a National Instruments NI-USB6259 data-acquisition system. Grounded channels are inserted between active measurement channels to reduce channel-crossover effects. Individual bank-capacitor voltages are collected for D1, whereas only the overall bank voltage can be acquired for D2 because the number of internal nodes exceeds the available channels. The backup-

Table 5.1: Experimental current range for each implemented configuration.

Topology	Stages	$I_{\min}$	$I_{\max}$
D1	3	90.000 µA	8250.000 µA
D2	3	90.000 µA	5250.000 µA
D1	4	90.000 µA	7650.000 µA
D2	4	90.000 µA	2250.000 µA
D1	5	90.000 µA	4650.000 µA
D2	5	90.000 µA	450.000 µA

capacitor voltage and the $\texttt{EN}_1$ and $\texttt{EN}_2$ signals are collected for both topologies. Texas Instruments TL081P low-input-bias buffers are used to isolate the system under test from the dynamic current drawn by the multiplexer while scanning its channels. An Agilent 3631A dual power supply provides the ±10 V supply required by the buffers.

The acquired bank-voltage traces are also used to identify the actual activation and deactivation instants of the source and sink SMUs around the $\texttt{EN}_1$ and $\texttt{EN}_2$ transitions. This accounts for the approximately 20 ms handover delay between source deactivation and sink activation. Input and output powers are reconstructed using the mean currents measured by the respective SMUs, with power set to zero during the handover interval. The corresponding energies are then obtained through trapezoidal integration.

Ultimately, the tested current range extends from the minimum current that permits cold start, 90 μA, to the maximum current at which each *CapDYN* configuration can switch stages reliably, as reported in Table 5.1. The upper limit depends on the smallest stage capacitance and on the 15 ms delay imposed by the LTC2934 voltage monitor. At higher charging currents, consecutive stage transitions can occur less than 15 ms apart, causing the monitor to miss them. A discrete implementation based on TLV3691 comparators and an ultra-low-power REF35 voltage reference could provide comparable power consumption while reducing the minimum detectable interval to the comparator response time, on the order of 40 μs. Since the resulting upper current limit differs among configurations, comparisons between D1 and D2 at the same stage count are restricted to their overlapping measured ranges.

The D1 and D2 banks were implemented using 220 µF capaci-

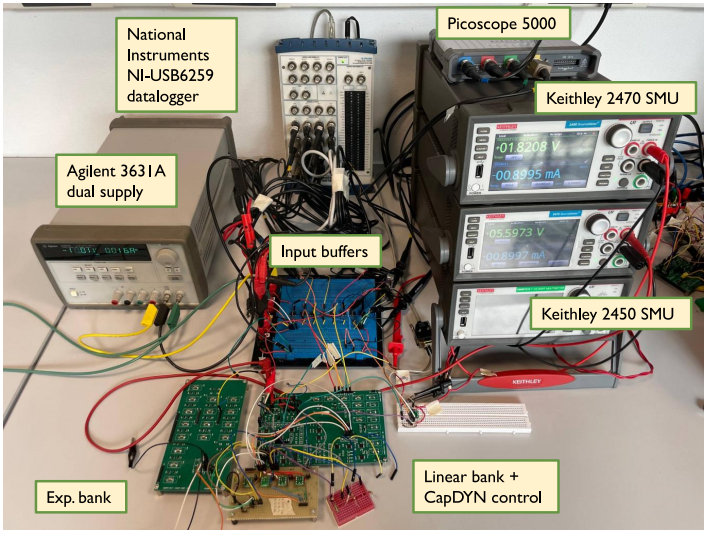


Figure 5.1: *CapDYN* efficiency measurement setup. On the top right, the source and sink SMUs, Keithley's 2470 and 2450, respectively. At the center, we can see National Instrument's NI-USB6259 data acquisition system. On the left, Agilent's 3631A dual power supply for the TL081P buffers (on the blue breadboard).

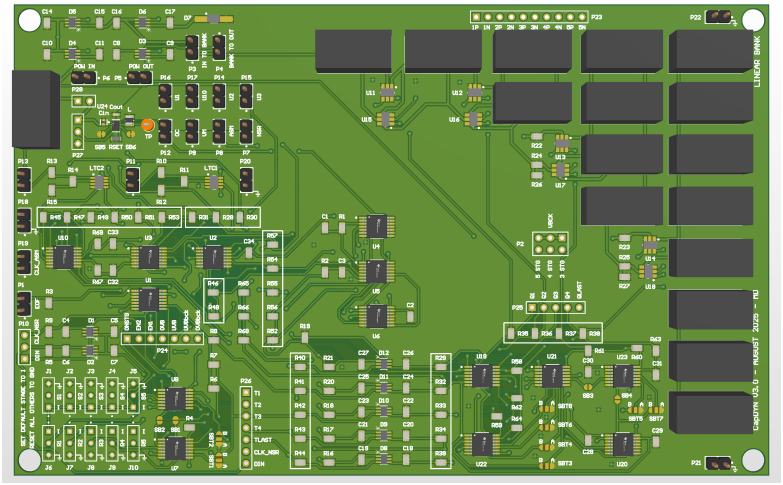


Figure 5.2: PCB of the CapDYN board, integrating the D1 bank design.

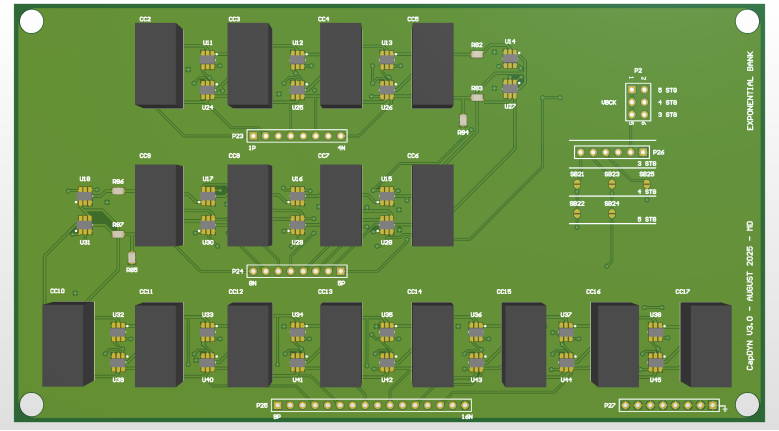


Figure 5.3: PCB of the D2 bank design.

tors (Murata, ceramic SMD, 6.3 V, 1206, ±20%, X5R, part number GRM31CR60J227ME11L), while the backup capacitor consisted of the parallel connection of one 220 µF capacitor and one 100 µF aluminum electrolytic capacitor (Jamicon, aluminum electrolytic radial, 35 V, 100 µF, ±20%, SK series, part number SKR101M1VE11M). The backup capacitance was selected so that the most power-consuming *CapDYN* configuration (5-stage, D2 bank) would not fully discharge the backup supply during the lowest-current efficiency test with 90 μA. The PCBs of the two banks and the control logic are shown in Fig. 5.4.

5.1.2 Hardware validation of stage switching and power-flow control

Table 5.2 summarizes the equivalent capacitance and available energy of each implemented *CapDYN* configuration. The reported energy corresponds to the minimum energy available in each stage of the considered configuration. In all cases, the final stage stores the largest share of capacitance and energy, as it corresponds to the configuration in which all capacitors are connected in parallel. Comparing the two topologies, D1 provides a higher equivalent capacitance in the early stages, whereas the last stages converge to the same equivalent capacitance. However, the available energy also depends on the post-switch voltage boost obtained during re-

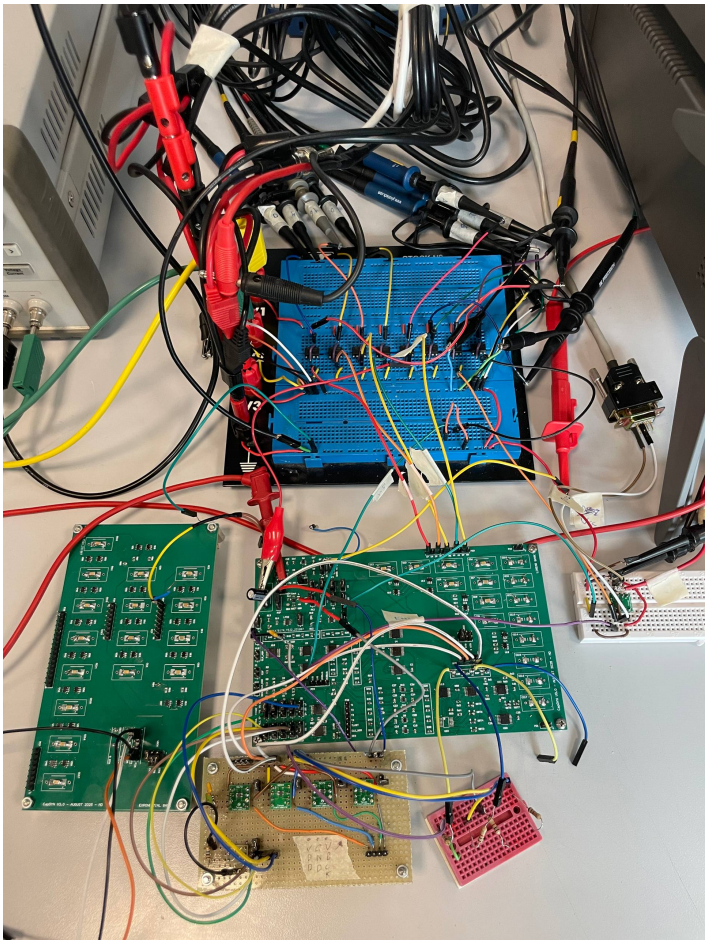


Figure 5.4: *CapDYN* prototype under development. The central PCB integrates the entire control logic, as well as the D1 bank. The PCB on the left integrates the D2 bank. The protoboard on the bottom includes the TLV3691 comparators used to interface progressive scaling mechanism and bank analog switches. At the top, the TL081P buffers, while on the right we see the $\texttt{EN}_2$ actuated switch to trigger the input/output SMUs.

configuration. As discussed in Appendix A.1, D2 produces a voltage doubling after each down-scale transition, while the voltage boost in D1 becomes progressively smaller as the system returns to lower stages. Therefore, although D2 has lower equivalent capacitance in the early stages, its post-switch voltage boost yields a slightly higher concentration of available energy across the stage sequence. Overall, D1 and D2 provide comparable stage-energy distributions despite their different capacitance profiles, as shown in Fig. 5.13 and Fig. 5.14. Consequently, their relative performance is expected to depend on two factors: the circuit-level cost of each topology, mainly switching losses and infrastructural consumption, and the degree to which the resulting stage-energy levels match the application task energies.

This second aspect is evaluated in the LTspice system-level study of Sec. 5.2, where the measured and modeled *CapDYN* behavior is coupled to an energy-admissible task scheduler. There, the stage-energy distribution becomes the scheduler-visible energy interface, allowing us to quantify how topology and task-to-stage matching affect throughput, responsiveness, and effective efficiency.

We first experimentally evaluate the charging and discharging behavior of the *CapDYN* hardware implementation and verify the operation of its automatic switching mechanism and I/O power-flow control for all implemented configurations, namely D1 and D2 banks with 3, 4, and 5 stages. For this functionality test, we charge and discharge the bank with a constant current of 90 μA. The operating voltage range, which also defines the switching thresholds as described in Sec. 4.2.2, is set to $[1.8\text{V}, 3.6\text{V}]$, consistently with TI MSP430FR-series microcontrollers, which are commonly adopted in intermittent systems. Since $V_H = 2V_L$, all D1 and D2 stages can be exploited without incurring brownout, as discussed in Sec. 4.2.2. A safety margin of 0.05 V is added to the upper threshold to prevent the switching mechanism from entering the *loop condition* due to the small equalization transients that characterize the D1 bank after some stage transitions.

Fig. 5.5 confirms that the implemented hardware follows the expected behavior derived from the capacitance and energy distributions in Tab. 5.2. During charging, *CapDYN* starts from stage 1, where the capacitors are arranged in the most series-oriented configuration and the bank exhibits its minimum equivalent capacitance. As the bank reaches the upper threshold, the automatic switching

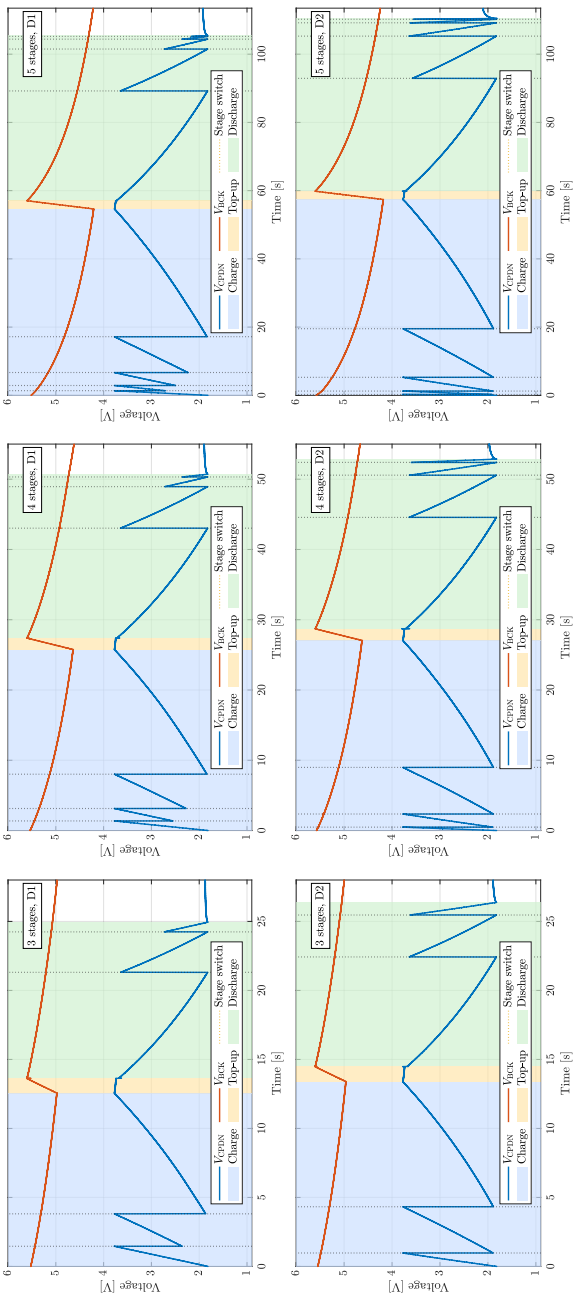


Figure 5.5: Charge/discharge voltage traces for the 90 μA reference test.

Table 5.2: Stage-wise equivalent capacitance and available energy.

(a) Available energy per stage (mJ)

N	Topology	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5
3	D2	0.267	1.337	5.773		
3	D1	0.178	1.247	5.684		
4	D2	0.134	0.668	2.807	11.679	
4	D1	0.089	0.446	2.584	11.457	
5	D2	0.067	0.334	1.403	5.680	23.425
5	D1	0.052	0.231	0.943	5.220	22.965

(b) Equivalent stage capacitance (mF)

N	Topology	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5
3	D2	0.055	0.220	0.880		
3	D1	0.088	0.220	0.880		
4	D2	0.028	0.110	0.440	1.760	
4	D1	0.080	0.176	0.440	1.760	
5	D2	0.014	0.055	0.220	0.880	3.520
5	D1	0.077	0.160	0.352	0.880	3.520

mechanism advances to the next stage, increasing the equivalent capacitance and allowing additional energy to be stored. This sequence continues until the final stage is reached, where the bank reaches its maximum equivalent capacitance and stored energy. The longer charge and discharge intervals observed as the number of stages increases are therefore consistent with the larger final-stage capacitance and energy availability reported in the table.

The traces also show the effect of the different capacitance profiles of D1 and D2. In the first stage, D2 charges faster than D1 because its minimum equivalent capacitance is smaller. Conversely, D2 reaches the end of the charging sequence later, since each capacitance scale-up is associated with a larger voltage decrease. This behavior is coherent with the slightly larger final-stage available energy reported for D2, even though the final-stage capacitance is the same for the two topologies.

The same measurements verify the correct operation of the I/O power-flow control. Once the target charging stage is reached, here the final stage, the control logic temporarily redirects the harvested

energy to top up the backup capacitor before enabling the load path. After this top-up interval, the discharge is sustained by *CapDYN*, while the backup capacitor supplies the auxiliary circuitry required to preserve the control functionality during task execution. Accordingly, the backup capacitor must be sized according to the longest application task to be supported, so that *CapDYN* can provide its voltage-scaling service throughout the full discharge interval.

During discharge, *CapDYN* progressively moves back toward lower stages. Whenever the voltage of the active stage drops below the lower threshold of 1.8 V, the capacitor bank is reconfigured to a lower stage. This transition rearranges the capacitors into a more series-oriented configuration and produces a post-switch voltage boost, allowing the system to continue supplying the load instead of leaving the remaining stored energy unusable below the operating voltage limit. This mechanism is at the core of the leftover energy minimization offered by *CapDYN* over a single capacitor (Sec. 5.1.4). The different boost behavior of the two topologies is clearly visible in the traces: D2 produces the expected voltage doubling after each down-scale transition, whereas the boost in D1 becomes progressively smaller as the bank returns to lower stages. When the bank is fully discharged down to the lower voltage threshold, at stage 1, the bank correctly interrupts the power flow to the load SMU, preserving the leftover energy and preventing uncontrolled depletion of the storage bank.

5.1.3 Efficiency metrics computation

By isolating charge and discharge instants, we can distinguish input and output energies:

$$E_{IN} = \int P_{IN}(t)\,dt \tag{5.1}$$

$$E_{OUT} = \int_{EN_2=H} P_{OUT}(t)\,dt \tag{5.2}$$

The input energy E_{IN} comprises the energy delivered to the main capacitor bank, $E_{\mathrm{IN,BANK}}$, and the energy routed through the backup path to establish and sustain the *CapDYN* control infrastructure, $E_{\mathrm{IN,BCK}}$. The $\mathtt{EN_1}$ signal directs the input power between these two paths: when $\mathtt{EN_1}$ is low, the backup capacitor is charged,

whereas when $\texttt{EN}_1$ is high, the input power is directed to the main bank. Accordingly,

$$E_{\mathrm{IN,BCK}} = \int_{EN_1=L} P_{\mathrm{IN}}(t)\, dt, \tag{5.3}$$

$$E_{\mathrm{IN,BANK}} = E_{\mathrm{IN}} - E_{\mathrm{IN,BCK}}. \tag{5.4}$$

We additionally distinguish between *cold-start* and *regimen* operation. Cold start is the initial interval required to charge a fully depleted bank to V_{SL} in stage 1. This condition defines the minimum viable bank state, as it marks the threshold beyond which additional stored energy becomes available to the load. Regimen operation begins from this state: a complete cycle includes charging from V_{SL} in stage 1 to V_{SH} in stage N, followed by discharge until the bank returns to V_{SL} in stage 1. During normal operation, to ensure the energy remains accessible to the load, task execution stops before the bank falls below the V_{SL} threshold in stage 1, keeping *CapDYN* within the regimen operating range. A new cold start is required only if the bank is depleted below this threshold, for example when the harvested input power is insufficient to sustain the infrastructural consumption during charging. Given the low infrastructural power reported in Fig. 5.6, *CapDYN* is expected to operate predominantly under regimen conditions.

The input energies associated with the regimen cycle are:

$$E_{IN,REG} = \int_{t\in REG} P_{IN}(t)\, dt \tag{5.5}$$

$$E_{\mathrm{IN,BCK,REG}} = \int_{t\in REG, EN_1=L} P_{\mathrm{IN}}(t) dt \tag{5.6}$$

$$E_{IN,BANK,REG} = E_{IN,REG} - E_{IN,BCK,REG} \tag{5.7}$$

The portion of E_{IN} that is necessary to overcome cold start but cannot be used to power the load is termed *leftover energy*, E_{LEFT}. We compute E_{LEFT} as the energy stored by the main bank when full load-driven discharge (V_{SL}, 1st stage) is achieved. E_{STORED} is computed either by summing the energies of each capacitor for D1 (Eq. 5.9) or based on the bank voltage for D2 (Eq. 5.10), as the number of bank capacitors is too large to be tracked by the available instrumentation for N $>$ 3.

$$E_{LEFT} = E_{\text{STORED}}(t = t_{EN2\downarrow}) \tag{5.8}$$

$$E_{STORED}^{D1} = \sum_{i=1}^{N} \frac{1}{2} C_i \left(V_{C_i,\text{pos}} - V_{C_i,\text{neg}}\right)^2 \tag{5.9}$$

$$E_{STORED}^{D2} = \frac{1}{2} C_{eq}^{D2}(j;N) V_{D2}^2 \tag{5.10}$$

Ultimately, $E_{\text{IN,BANK,REG}}$ is the maximum energy available for delivery to the load. The portion not delivered as E_{OUT} defines the aggregate bank-side loss E_{RES} and accounts for the bank's parasitic effects:

$$E_{RES} = E_{IN,BANK,REG} - E_{OUT} \tag{5.11}$$

This residual includes conduction, leakage, switching losses, and other bank-side parasitic contributions, which are not individually separated by this metric.

Having defined all the energies at play, we now focus on defining the efficiency metrics useful for comparing *CapDYN* against the single capacitor benchmark. The efficiency η_{CS} compares total input and output energies when cold starting the system.

$$\eta_{CS} = \frac{E_{OUT}}{E_{IN}} \tag{5.12}$$

$\eta_{CS,BANK}$ compares bank input and output energies when cold starting the system. This metric represents the upper boundary of η_{CS}, if infrastructural power consumption tends to zero.

$$\eta_{CS,BANK} = \frac{E_{OUT}}{E_{IN,BANK}} \tag{5.13}$$

The efficiency η_{REG} compares regimen input and output energies.

$$\eta_{REG} = \frac{E_{OUT}}{E_{IN,REG}} \tag{5.14}$$

The efficiency $\eta_{REG,BANK}$ compares bank input and output energies in regimen. This metric represents the upper boundary of η_{REG}, if infrastructural power consumption tends to zero.

$$\eta_{REG,BANK} = \frac{E_{OUT}}{E_{IN,BANK,REG}} \tag{5.15}$$

We then express infrastructural consumption, leftover energy, and aggregate bank-side loss as normalized input-energy fractions. A metric without a boundary suffix is normalized by the complete-system input energy E_{IN}; the suffix BANK denotes normalization by $E_{\mathrm{IN,BANK}}$, while BANK, REG denotes normalization by the bank input energy during regimen operation, $E_{\mathrm{IN,BANK,REG}}$.

$$\lambda_{INFR} = \frac{E_{IN,BCK}}{E_{IN}} \tag{5.16}$$

$$\lambda_{LEFT} = \frac{E_{LEFT}}{E_{IN}} \tag{5.17}$$

$$\lambda_{RES} = \frac{E_{RES}}{E_{IN}} \tag{5.18}$$

The following test results respond to where a reconfigurable capacitance structure like *CapDYN* can improve fixed capacitor supplies in terms of the aforementioned metrics. Indeed, the benchmark is a single capacitor which is able to supply the same amount of output energy to its load, with respect to *CapDYN*. In these tests, we assume that the capacitor charges up to V_{SH} and then discharges down to V_{SL}: at that voltage, the single capacitor cannot supply power to the load anymore, so the remaining energy is considered leftover energy. Although the single capacitor cannot reconfigure itself or ensure energy-admissible task execution as *CapDYN*, the single capacitor features neither infrastructural power consumption nor switching losses. Thus, its charging and discharging efficiency is mostly dictated by the capacitor's own series resistance, which we assume negligible at the tested current levels.

5.1.4 Infrastructural power, efficiencies, and losses as a function of current

Current infrastructural power-consumption values are shown in Fig. 5.6. These results were obtained through a Keithley 2470 SMU by sweeping the backup-capacitor voltage from 5.5 V down to 3 V in 15 steps, each lasting 20 s, while measuring the drawn current. Each reported value is the mean power measured during the corresponding voltage step. As expected, power consumption grows with backup voltage: the minimum value, 4.15 µW, is achieved by the 3-stage D1 configuration at 3 V, whereas the maximum value, 9.07 µW, is reached

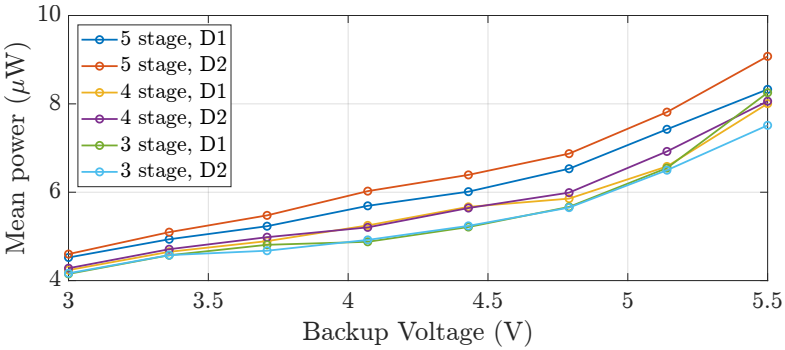


Figure 5.6: Power consumption of CapDYN with 3, 4, 5 stages, matched with either D1 or D2 bank.

by the 5-stage D2 configuration at 5.5 V. At three stages, the power consumption of D1 and D2 is similar, whereas the difference between the two topologies grows with stage number. This trend is consistent with the poorer scalability of D2, whose exponential increase in switch count causes a steeper rise in infrastructural power consumption.

Although the entire infrastructure consumes less than 10 µW even in the worst case, the corresponding energy overhead can still represent a large fraction of the harvested energy, as shown in Fig. 5.7. Cold start requires a minimum input current of 90 µA to activate the buck converter. At such low charging and discharging current levels, the backup capacitor remains in discharge for longer because it must sustain infrastructural power consumption over a longer interval. In addition, the first charging event of the backup capacitor during cold start features transients that delay the onset of effective charging. Together, these effects extend the cold-start interval and increase the amount of energy absorbed by the infrastructure at low currents.

More generally, infrastructural energy is expected to decrease as current increases, since a faster charge/discharge trajectory shortens the time over which infrastructural power is drawn. Approximating infrastructural power as nearly constant,

$$E_{IN,\mathrm{BACKUP}} \approx P_{\mathrm{INFR}}\, T \propto \frac{1}{I}, \tag{5.19}$$

because the charge/discharge duration T is inversely proportional

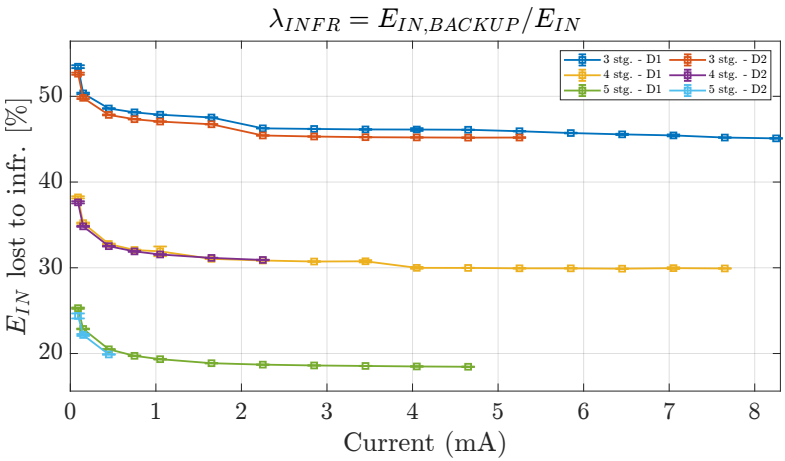


Figure 5.7: Percentage of input energy lost to infrastructural consumption.

to current. Thus, $E_{IN,\text{BACKUP}}$ is expected to exhibit an approximately hyperbolic dependence on current. The normalized metric λ_{INFR} does not retain a purely hyperbolic profile, but it still inherits part of this behavior because $E_{IN,\text{BACKUP}}$ is a component of E_{IN}. Since $E_{IN,\text{BACKUP}}$ varies more strongly with current than the total input energy, the ratio $\lambda_{INFR} = E_{IN,\text{BACKUP}}/E_{IN}$ preserves hyperbolic features, which explain its increase in the low-current region of λ_{INFR} in Fig. 5.7 below 450 µA. Above 450 µA, cold start proceeds without the additional low-current transients, and λ_{INFR} transitions toward the smoother baseline trend set by the approximately hyperbolic dependence of $E_{IN,\text{BACKUP}}$ on current.

Two downward folds are also visible: at 1650 µA for both 3-stage configurations, and at 3450 µA for the 4-stage D1 configuration. Above these currents, the corresponding configurations no longer execute backup top-up before entering discharge, and therefore exhibit a slightly smaller λ_{INFR}. This occurs because the backup capacitor does not discharge sufficiently during charging to trip the voltage-monitor hysteresis and trigger a top-up event Sec. 4.2.3. The same folds are visible again in Fig. 5.10a. Finally, as the number of stages increases, λ_{INFR} decreases. Although infrastructural power consumption, and thus $E_{IN,\text{BACKUP}}$, grows slightly with stage number, as shown in Fig. 5.6, the last-stage energy capacity, and therefore the

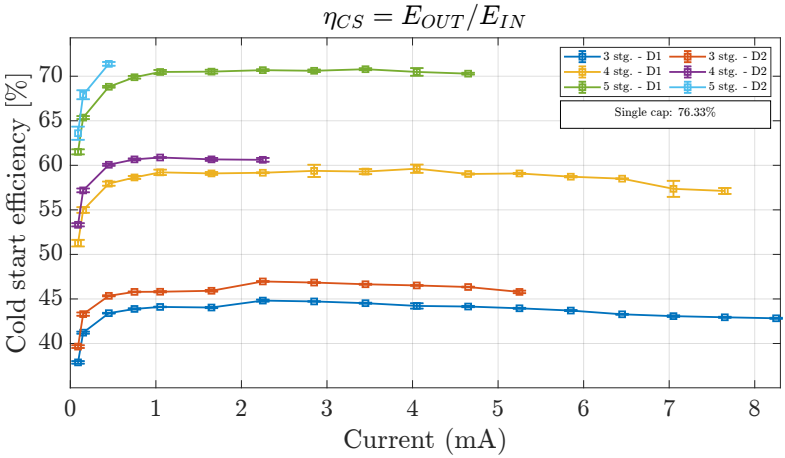
overall input-energy budget E_{IN}, grows much faster. Consequently, the fraction of input energy spent on infrastructure becomes progressively smaller, resulting in an appreciable decrease of λ_{INFR} for each additional stage.

Given the strong impact of infrastructural consumption on *CapDYN* efficiency, we distinguish between system-level efficiencies, namely η_{CS} and η_{REG}, which include the entire energy budget flowing into and out of the platform, and bank efficiencies, namely $\eta_{CS,\text{BANK}}$ and $\eta_{REG,\text{BANK}}$, which isolate the bank contribution. The latter can be interpreted as the upper bound achievable by the selected topology in the limit case of negligible infrastructural consumption.

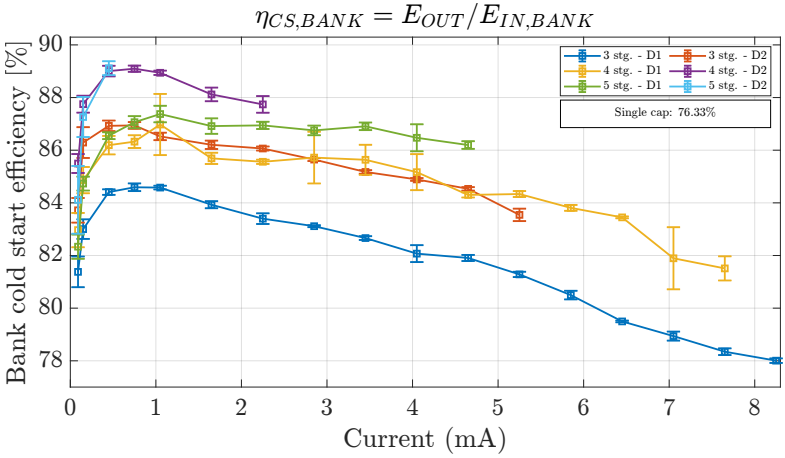
Cold-start efficiency results are shown in Fig. 5.8. The system-level metric η_{CS} in Fig. 5.8a is strongly influenced by infrastructural consumption. As the number of stages increases, η_{CS} rises while λ_{INFR} decreases, because infrastructural energy grows more slowly than both the total input energy and the useful output energy. Consequently, a smaller fraction of the harvested energy is dissipated in the infrastructure and a larger fraction is delivered to the load. The best cold-start efficiency is achieved by the 5-stage D2 configuration, which reaches 71.4%. This value approaches the efficiency of a single capacitor (76.33%), while retaining the task-adaptive benefits of *CapDYN*, whose application-level advantages will become clearer in the LTspice evaluation section, Sec. 5.2

At bank level, Fig. 5.8b shows that once infrastructural consumption is removed, the single-capacitor benchmark can be clearly exceeded. This improvement is mainly enabled by the voltage-boosting behavior of *CapDYN*, which strongly reduces leftover energy at full discharge. Peak bank cold-start efficiency is achieved in the range 450 µA–1050 µA, with a maximum value of 89.1% for the 5-stage D2 configuration. Beyond this region, efficiency tapers because resistive-path losses progressively increase with current, as quantified in Fig. 5.11c. Four- and five-stage configurations perform similarly, whereas three-stage configurations are consistently less efficient, with 3-stage D1 giving the lowest bank cold-start efficiency, although still above the single-capacitor benchmark. The origin of this trend becomes clearer when considering leftover-energy losses.

The reason why bank cold-start efficiency exceeds the single-capacitor benchmark can be found in the $\lambda_{LEFT,\text{BANK}}$ results shown in Fig. 5.9b. A larger leftover energy directly reduces the usable out-

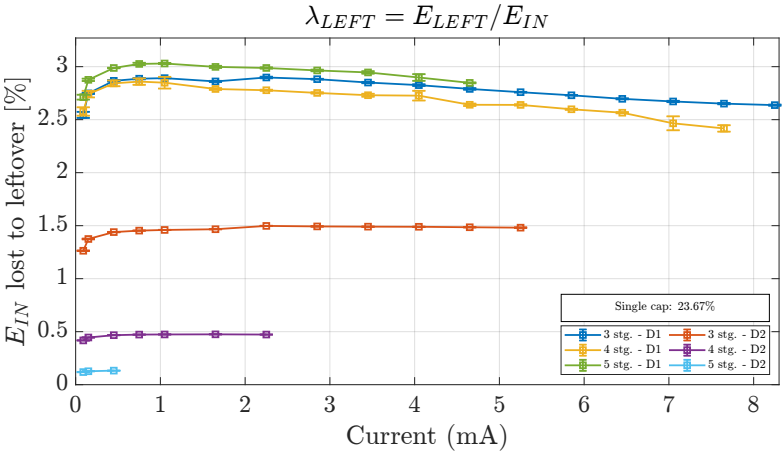


(a) Cold start energy efficiency of *CapDYN* as a function of charge/discharge current, in comparison to a single capacitor benchmark.

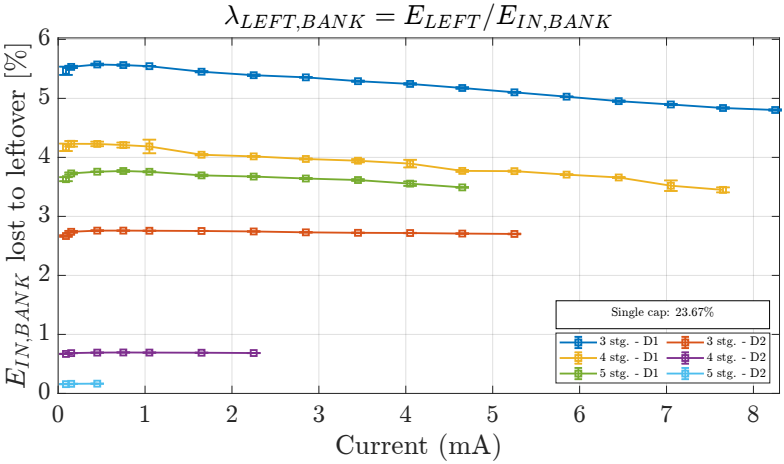


(b) Cold start energy efficiency of the *CapDYN* capacitor bank, as a function of charge/discharge current.

Figure 5.8: Cold start efficiency metrics of *CapDYN* as a function of charge/discharge current. At the top, the entire infrastructure's cold start efficiency is displayed, while at the bottom we focus only on the bank's cold start efficiency.



(a) Percentage of input energy lost to leftover energy in *CapDYN*, as a function of charge/discharge current.



(b) Percentage of input energy lost to leftover energy in the *CapDYN* capacitor bank, as a function of charge/discharge current.

Figure 5.9: Percentage of input energy lost to leftover energy as a function of charge/discharge current. The top panel includes the full system, while the bottom panel focuses on the bank only.

put energy and therefore lowers $\eta_{CS,\mathrm{BANK}}$. Four- and five-stage configurations, for both D1 and D2, exhibit similarly low $\lambda_{LEFT,\mathrm{BANK}}$, whereas the 3-stage configurations perform worse, with 3-stage D1 representing the least favorable case. This trend follows the first-stage capacitance of each configuration (Tab. 5.2): the smaller the first-stage capacitance, the smaller the residual leftover energy, and thus the higher the bank efficiency.

More importantly, $\lambda_{LEFT,\mathrm{BANK}}$ captures the actual benefit of *CapDYN* over a single-capacitor benchmark. While a single capacitor leaves 23.67% of the input energy trapped as leftover energy, *CapDYN* reduces this fraction to as little as 0.16% in the best case, namely 5-stage D2. Even in the least favorable case, 3-stage D1, leftover energy still accounts for only 5.57% of the bank input energy. This marked reduction is one of the main advantages of *CapDYN* and directly contributes to the improvement of η_{CS}. This reduction is important not only for energy utilization but also for responsiveness: the same capacitance-scaling mechanism that reduces unusable residual energy also shortens the time required to return to an executable energy state.

For completeness, the corresponding system-level metric λ_{LEFT} is also shown in Fig. 5.9a. The closer spacing of the D1 curves in λ_{LEFT} arises because

$$\lambda_{LEFT} = \lambda_{LEFT,\mathrm{BANK}}(1 - \lambda_{INFR}).$$

Although D1 exhibits larger stage-dependent differences in $\lambda_{LEFT,\mathrm{BANK}}$, these are partially compressed by λ_{INFR}, whereas the intrinsically wider separation of the D2 $\lambda_{LEFT,\mathrm{BANK}}$ curves remains visible even after scaling by $(1 - \lambda_{INFR})$.

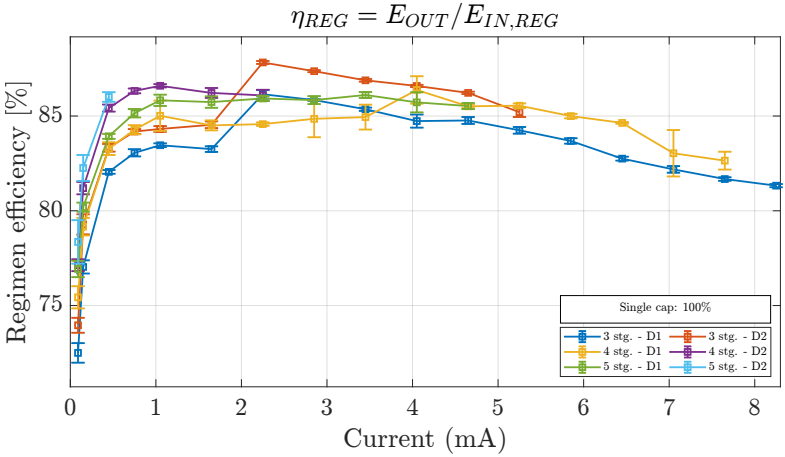
An interesting measurement-accuracy consideration also emerges from $\lambda_{LEFT,\mathrm{BANK}}$. This quantity slightly decreases with increasing current. When leftover energy is estimated from the bank terminal voltage, higher current increases the resistive voltage drop during both charge and discharge. As a result, the leftover energy is progressively underestimated, leading to a smaller estimated ratio $E_{LEFT}/E_{IN,BANK}$ at higher current. Allowing a suitable relaxation time before measuring E_{LEFT} would likely mitigate this bias by reducing the effect of the resistive drop. Owing to setup limitations, however, this verification was left for future work, so the low-current measurements should be regarded as the least affected by this source of bias.

Estimating energy at the system boundary, instead, avoids the current-dependent bias associated with terminal-voltage-based stored-energy estimation. At the same time, the observed current dependence is not merely a measurement artifact: under a fixed terminal-voltage cutoff, higher current genuinely increases ohmic losses and reduces the effective internal voltage range, since the cutoff is reached at a lower internal state of charge. Therefore, the corresponding input- and output-energy measurements, performed at the system boundary, capture the actual exchanged energy and can be regarded as representative of the true system efficiency.

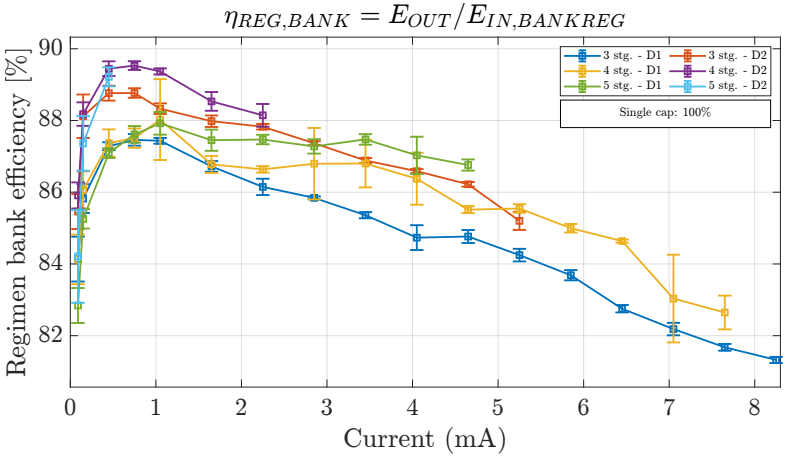
The cold-start discussion also provides the basis for interpreting regimen efficiency: once infrastructural effects are separated out and leftover-energy behavior is understood, the remaining differences between D1 and D2 can be attributed mainly to the balance between resistive-path and switching-related losses.

Regimen bank efficiency results in Fig. 5.10b show that D2 is more efficient than D1 at low current, while D1 progressively improves relative to D2 as current increases, especially at larger stage counts. This behavior stems from the different loss mechanisms that dominate the two topologies. D2 has a slightly longer resistive path at 4 and 5 stages, due to its larger number of switches, but it also benefits from lower switching losses because it nominally avoids equalization transients. At low current, time-dependent parasitic effects become more relevant than resistive-path losses, favoring D2 in the measured comparisons. At higher current, however, conduction losses along the longer D2 path become increasingly significant, reducing its advantage and allowing D1 to recover. The current sensitivity of each topology becomes weaker as stage number increases, because the larger bank energy capacity of higher-stage systems reduces the relative impact of resistive-path loss. In the 3-stage case, D2 and D1 exhibit comparable slopes because the switch count of 3-stage D2 is still limited. Similarly to λ_{LEFT} and $\lambda_{LEFT,\mathrm{BANK}}$ curves, $\eta_{REG,BANK}$ is η_{REG} scaled by $(1-\lambda_{INFR})$: η_{REG} is shown in Fig. 5.10a for completeness.

The role of bank-side losses is summarized in Fig. 5.11. Bank-side losses are defined by Eq. 5.11 and represent the fraction of regimen bank energy that is dissipated parasitically rather than delivered to the load as E_{OUT}. The relevant mechanisms are resistive-path losses, switching and equalization losses, and capacitor/switch leakage. Their different current dependences help explain the over-

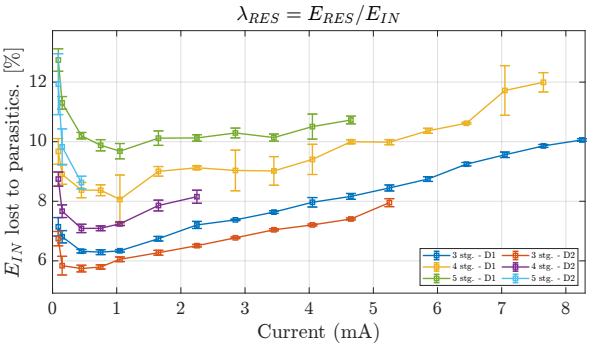


(a) Regimen energy efficiency of *CapDYN*, as a function of charge/discharge current.

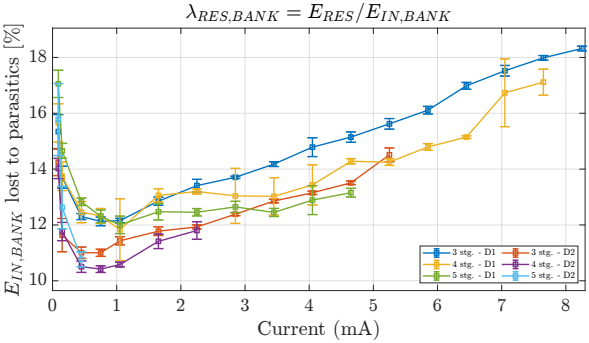


(b) Regimen energy efficiency of the *CapDYN* capacitor bank, as a function of charge/discharge current.

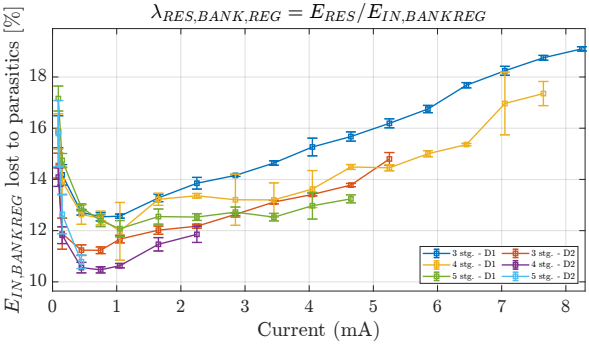
Figure 5.10: Regimen efficiency of *CapDYN* as a function of charge/discharge current. At the top, the entire infrastructure's regimen efficiency is displayed, while at the bottom we focus only on the bank's regimen efficiency.



(a) Percentage of cold-start input energy lost to parasitic losses in *CapDYN*, as a function of charge/discharge current.



(b) Percentage of cold-start input energy lost to parasitic losses in the *CapDYN* capacitor bank, as a function of charge/discharge current.



(c) Percentage of regimen input energy lost to parasitic losses in the *CapDYN* capacitor bank, as a function of charge/discharge current.

Figure 5.11: Parasitic input energy losses versus charge/discharge current, normalized to total cold-start, bank cold-start, and bank regimen input energy.

all shape of Fig. 5.11c: the measured trends are consistent with conduction-related losses at high current and time-dependent parasitic effects at low current, while switching and equalization losses depend on the topology-dependent voltage mismatch at the transition instants. Accordingly, the sub-mA peak in bank-side losses, which is responsible for the dip observed in all bank-efficiency curves, is consistent with a time-dependent contribution, while resistive-path effects are expected to be less influential in that region.

Resistive-path losses affect both $E_{IN,\mathrm{BANK,REG}}$ and E_{OUT}: as current increases, they raise the required input energy and reduce the useful output energy. Therefore, they are expected to produce a monotonic increase of

$$\lambda_{RES,\mathrm{BANK,REG}} = \frac{E_{IN,\mathrm{BANK,REG}} - E_{OUT}}{E_{IN,\mathrm{BANK,REG}}}$$

with current. This interpretation is consistent with Fig. 5.11c, where the normalized regimen bank-loss term increases toward the high-current region.

Leakage, by contrast, is expected to produce the opposite trend. As current decreases, the charge/discharge trajectory becomes slower, and the bank spends longer at elevated voltage, where leakage can deplete stored energy. If leakage current is approximately independent of bank voltage, the resulting leakage energy should scale mainly with the charge/discharge duration and therefore follow an approximately hyperbolic, i.e., inverse-current, dependence, similarly to $E_{IN,\mathrm{BACKUP}}$ in Eq. 5.19. If leakage instead increases with voltage, low-current operation should be even more strongly affected, since the bank remains longer in the high-voltage region where leakage is larger. In all cases, however, leakage is expected to produce a smooth monotonic, hyperbola-like trend with current.

Switching and equalization losses only affect $E_{IN,\mathrm{BANK}}$, because stage-decrease operations during discharge do not involve equalization transients, whereas stage-increase operations during charge do. To estimate whether D1 equalization losses could contribute to the sub-mA loss peak, we measured the voltage mismatch between the capacitor groups that become parallelized during each D1 charge-stage transition and applied the ideal charge-sharing loss expression of Eq. (A.7), derived in Appendix A.1. The corresponding total

switching-loss energy was estimated as

$$E_{SW,\mathrm{LOSS,TOT}} = \sum_{k=1}^{N-1} \frac{1}{2} \frac{C_{A,k} C_{B,k}}{C_{A,k} + C_{B,k}} (\Delta V_k)^2 \tag{5.20}$$

where $C_{A,k}$ and $C_{B,k}$ are the equivalent capacitances of the two groups connected in parallel at transition k, and ΔV_k is their measured pre-switch voltage mismatch. D2 switching losses could not be estimated in the same way because of the larger number of involved capacitors and internal nodes. However, switching losses are not expected to be the dominant loss mechanism for D2. Indeed, according to the ideal charge-sharing loss expression of Eq. (A.7), no redistribution loss occurs when the reconfigured capacitor groups have equal voltage. This is the case for the ideal D2 transitions, which connect or split equal-voltage groups, as discussed in Appendix A.1.

As shown in Fig. 5.12, the estimated D1 switching-loss term exhibits a mild low-current rise, consistent with the sub-mA feature observed in Fig. 5.11c. This indicates that switching losses do contribute to the low-current loss peak, although only marginally. At intermediate currents, switching losses reach a plateau whose width decreases as the stage number increases, and then fall again at higher currents.

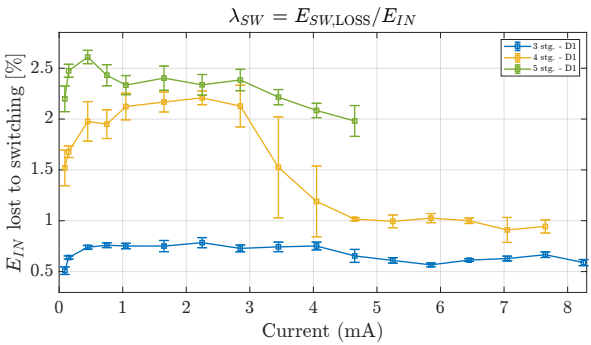
This behavior should be interpreted in terms of the voltage mismatch that exists, immediately before each switching event, between the capacitors that are about to be connected in parallel. The equalization loss is set by these pre-switch voltage differences: larger mismatches produce larger charge-sharing transients, whereas smaller mismatches reduce the switching-loss contribution. Therefore, the current-dependent trend observed in Fig. 5.12 should not be attributed to the switch event alone, but to how the charge/discharge current affects the pre-switch voltage distribution across the capacitors before reconfiguration. Since mid-range currents produce the highest estimated equalization loss, the measured trend implies that the pre-switch voltage mismatch between capacitor groups is also maximized in this current range. The mechanism causing this non-monotonic mismatch is not yet clear from the available measurements. Identifying the detailed origin of this current-dependent mismatch is left to future work.

In any case, the estimated D1 equalization loss accounts for only a limited portion of the sub-mA efficiency degradation. A compa-

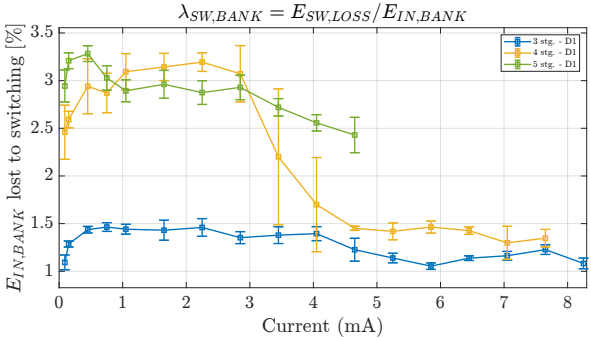
rable feature is also observed for D2 in Fig. 5.11c, although its ideal transitions connect equal-voltage groups and nominally avoid redistribution loss. Equalization therefore cannot explain the sub-mA feature shared by both topologies. Its occurrence at low current instead points to a time-dependent parasitic contribution, since the longer charge–discharge interval of the low-current tests allows more energy to be depleted through capacitor, switch, or board leakage. Conversely, the high-current trend is consistent with resistive-path effects. The individual contributions were not independently separated by the present measurements.

Overall, this current-based characterization highlights both the strengths and the limitations of *CapDYN*. Its main advantage over a single capacitor is the strong reduction of leftover energy, which allows the bank-level cold-start efficiency to surpass the single-capacitor benchmark while also exposing intermediate energy states for energy-admissible task execution. Its main drawbacks are the infrastructural overhead required to monitor and reconfigure the bank, and the additional bank-side parasitic losses associated with resistive paths, switching activity, and leakage, which become especially relevant in the low-current regime.

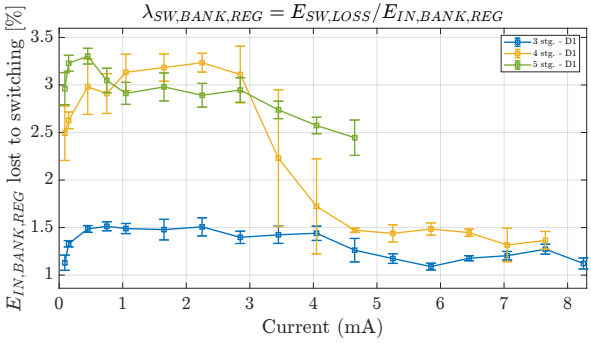
This characterization also clarifies why hardware efficiency alone does not fully describe the benefit of *CapDYN*. The hardware tests intentionally exercise the storage subsystem as an isolated energy buffer: the bank is charged to the final stage and then fully discharged. This test mode is useful to quantify losses, leftover energy, and infrastructural consumption under controlled conditions, but it does not correspond to the intended use of *CapDYN* in a battery-free application. In the intended operating mode, the load does not need to wait for the final stage. Instead, the *CapDYN* energy-state interface allows the runtime to execute tasks as soon as the corresponding energy-admissible stage is reached, rather than waiting for complete charge and subsequent full discharge. The LTspice evaluation therefore builds on the hardware-tested stage structure, power consumption, and switching behavior to quantify how these mechanisms translate into application-level benefits when the intermediate energy states are actually exploited. As shown in the following section, coupling *CapDYN* to an energy-admissible scheduler allows its benefits to emerge more clearly, particularly in terms of successful throughput, first-completion time, timely yield, and effective efficiency.



(a) Percentage of cold-start input energy lost to switching equalization transients in *CapDYN*, as a function of charge/discharge current.



(b) Percentage of cold-start input energy lost to switching equalization transients in the *CapDYN* capacitor bank, as a function of charge/discharge current.



(c) Percentage of regimen input energy lost to switching equalization transients in the *CapDYN* capacitor bank, as a function of charge/discharge current.

Figure 5.12: Switching equalization transient losses versus charge/discharge current for the D1 bank, normalized to: total cold-start input energy, bank cold-start input energy, and bank regimen input energy.

5.1.5 Software and hardware controlled switching

In order to estimate what strategy is more efficient at controlling the previously introduced capacitor banks, we evaluate the benefits of *CapDYN*'s autonomous capacitance scaling and I/O control with respect to a software-based *CapDYN* version. In the software-based version, we replace the I/O power flow (Sec. 4.2.3) and capacitance scaling (Sec 4.2.2) section with an MSP430FR5994. The autonomous control section introduces an overhead of 1.036 µW, while the equivalent software-based control adds 2.38 µW and 3.65 µW of additional overhead, considering a monitoring frequency of 1 Hz and 50 Hz, respectively. The overhead introduced by the microcontroller is computed considering the MCU's low power mode power consumption, ADC measurements cost, and active mode cost during the stage monitoring. The results indicate that the autonomous *CapDYN* consumes less power while providing a higher reactivity compared to the software-based *CapDYN*. Although the overhead of the MCU-based model could be reduced by relaxing the monitoring frequency, which can be a solution for larger storages/higher-power scenarios, this reduces the system's reactivity to its active periods, which is problematic for smaller storages/lower-power scenarios, where more frequent charge monitoring is needed. *CapDYN* solves this reactivity-efficiency tradeoff, as it is always instantly reactive, with a fixed ultra-low power overhead, extending the application of dynamic capacitance to ultra-low power energy sources.

A noteworthy software-control solution is presented by Morphy, which can be identified as *CapDYN*'s closest work, as it shares the progressive capacitance scaling approach with *CapDYN*. To further motivate our autonomous control choice, we compare *CapDYN*'s constant overhead to Morphy. Morphy's average power consumption increases as the capacitor size scales down, due to a faster monitoring frequency. Indeed, a 4-stage Morphy with C = 100 mF, 10 mF, and 1 mF capacitors consumes 2.44 µW, 6.44 µW, and 46.99 µW, with 5, 0.5, and 0.05 s monitoring periods, respectively. *CapDYN*'s 4-stage implementation consumes a maximum of 8.06 µW for the D2 topology, regardless of the capacitor size (Sec. 5.1.4). Morphy achieves a lower power consumption for larger capacitances by reducing its monitoring period, assuming that changes are slower. This reduces its reactivity to voltage drops and input power bursts.

Therefore, Morphy might not guarantee the prevention of power failures at lower stages if its duty cycle is not properly adjusted. On the contrary, *CapDYN* is always reactive, which is especially important in the case of smaller capacitors: for the 1 mF case, *CapDYN* consumes up to 83% less average power than Morphy, besides requiring no software overhead for bank control. If the 1 mF Morphy were to achieve instant reactivity as *CapDYN*, Morphy would need to keep its control MCU always-on, which would lead to a power consumption 1000 times larger than *CapDYN*.

Besides controlling its progressive scaling with an MCU, Morphy also implements a software-based pre-regulation of its output voltage. Morphy separates a set of its bank (task capacitors) to power a certain task. Said task capacitors are maintained within an user-programmed range by charge-pump boosting, where the task capacitors are briefly connected to the remainder of the bank, which sits at a higher voltage. Although this approach allows more flexibility in terms of load voltage range than *CapDYN*'s output voltage limitation to $[V_{SL}, V_{SH}]$, Morphy's core/task voltage boosting solution is entirely based on lossy charge transfers. For example, if a 1 mF bank is charged at 3.6 V and is briefly connected to a 1 mF task capacitor, which is nearing the load's lower voltage limit of 1.8 V, 10% of the total energy is lost due to resistive dissipations in the charge transfers. If we factor in the overhead of a software-based control, this effect highly reduces Morphy's efficiency. We believe charge transfers should be limited by design, so we do not consider this as a viable regulation strategy.

5.2 Application-level LTspice evaluation

The previous hardware evaluation validates the operation and loss mechanisms of the *CapDYN* storage subsystem in isolation. We now evaluate how those mechanisms affect a complete intermittent application. For this purpose, LTspice is used to extend the hardware characterization to application-level design-space exploration. The simulated storage subsystem acts as a *CapDYN* digital twin, implementing the same capacitance-reconfiguration principles and input/output power-flow control verified in the prototype while allowing topology, stage count, and capacitance values to be varied systematically. Coupling this model to the application workload

and energy-admissible scheduler allows us to determine how different design choices affect system-level performance under realistic battery-free operating conditions, including harvested-power input traces and profiled task execution.

In particular, the model extends the design space beyond the 3-, 4-, and 5-stage bank sizes implemented in hardware to 6- and 7-stage configurations, which are evaluated only in LTspice because of their component count, PCB area, and cost, but are useful to assess whether additional energy states improve application-level performance enough to justify their overhead.

This section therefore evaluates *CapDYN* at application level by combining three elements: a profiled battery-free sensing workload, realistic harvested-power traces, and the LTspice digital twin of the storage subsystem. We first build and profile a load prototype executing the target sensing application, and use the measured task energies to derive the task-to-stage mappings used to dimension each *CapDYN* configuration. We then simulate the complete task cycle under realistic harvested-energy input, comparing D1 and D2 banks across different stage counts and evaluating how the energy-state interface affects task completion, responsiveness, timely yield, and effective efficiency. This evaluation allows us to assess both the practical feasibility of energy-admissible execution under intermittent power and the design-space tradeoffs that must be considered before integrating *CapDYN* into a final battery-free demonstrator.

5.2.1 Target battery-free sensing workload

The target load is a self-powered environmental sensing platform representative of precision-agriculture deployments in remote or infrastructureless locations. The application periodically samples temperature, humidity, and light intensity, captures and compresses images for visual inspection, transmits data through LoRa, and performs I/V sweeps to characterize the electrical impedance of soil or plant tissue, enabling non-invasive monitoring of environmental and crop health parameters [119, 120, 25].

The platform is built around a Texas Instruments MSP430FR-series microcontroller, selected for its ultra-low-power operation and integrated non-volatile FRAM. The sensing interface includes an SHT41 temperature and humidity sensor, a TSL2591 light sensor, and an HM0360 low-power image sensor, while wireless communica-

tion is provided by an SX1262 LoRa transceiver. The platform operates exclusively from harvested energy stored in *CapDYN*, without batteries. A lightweight runtime running on the microcontroller performs task scheduling, energy-admissibility checks, and coordination with the energy-state interface introduced in Sec. 4.3. Tasks are executed only when the current storage stage satisfies the admissibility condition of Sec. 4.3.1.

Task-level energy consumption is characterized offline through repeated executions while monitoring current with a precision shunt resistor and a digital oscilloscope. The resulting task energy profiles are used to derive the task-to-stage mapping of Sec. 4.4.4, which is later integrated in the LTspice simulations used for long-term evaluation under harvested-power traces.

The workload, summarized in Table 5.3, is designed to stress the execution model with heterogeneous tasks typical of autonomous batteryless sensing nodes. It includes low-power, short-duration sensing operations, such as temperature, humidity, and ambient light acquisition (T_1), which are typically executed frequently and benefit from rapid energy accumulation and low startup latency. More energy-demanding operations are represented by impedance-based measurements (T_2), which characterize soil or plant conditions through I/V sweeps. This task exhibits long execution times and sustained power consumption, making it particularly sensitive to energy availability and unsuitable for speculative execution. Its long duration is due to the low scan rate required to obtain meaningful impedance measurements of soil or plant tissues [9].

The transmission of the data acquired in T_1 and T_2 is modeled by T_3, whose energy cost includes radio activation and LoRa transmission overhead. Applications requiring retransmissions or redundancy can account for them by profiling the corresponding transmission sequence as a larger task or by scheduling repeated transmission tasks. The task set further includes image acquisition through the BW camera (T_4) and local image compression (T_5), which model bursty workloads combining short high-power operations with medium-duration processing phases. Finally, compressed image transmission (T_6) represents one of the most energy-intensive operations in the application, requiring prolonged radio activity and sustained power availability to complete successfully.

Task ordering is determined primarily by the functional dependencies of the application. Environmental and impedance data must

Table 5.3: Description of the implemented tasks

Tasks ID	Description	Class
T_1	Environmental Sensing	P0
T_2	I/V Sweep - Soil/Plant Assessment	P2
T_3	I/V Sweep and Env. Data Transmission	P2
T_4	Image Capture	P1
T_5	Jpeg Compression	P1
T_6	Image Data Transmission	P2

be acquired before they can be transmitted, while image data must be captured and compressed before image transmission. Therefore, the workload follows a sensing–processing–communication pipeline, rather than an arbitrary order selected to favor the proposed architecture. When multiple tasks do not have direct dependencies, their execution order can be selected by the scheduler or by the application logic; however, dependent operations must preserve the partial order imposed by the application.

Overall, the selected workload combines three classes of tasks:

1. **(P0)** - Frequent, low-energy sensing tasks;
2. **(P1)** - Medium-duration processing and data handling tasks;
3. **(P2)** - High-energy, long-running communication and measurement tasks.

This heterogeneity provides a meaningful benchmark for evaluating how energy-admissible execution coordinates computation with dynamically available energy while avoiding partial executions under intermittent power supply. It also represents a challenging condition for fixed-capacitor systems, because low-energy and high-energy tasks coexist in the same application flow, stressing the mismatch between fixed storage capacity and task-level energy demand. The evaluation therefore focuses on the task-energy distribution and task-to-stage matching that determine the behavior of reconfigurable storage in battery-free sense-and-send applications, rather than on arbitrarily large general-purpose task graphs.

5.2.2 Evaluation methodology

We evaluate the proposed architecture using real-world prototyping for the application load and task scheduler, and LTSpice simulations for the *CapDYN* storage hardware. The evaluation considers

Table 5.4: Task power, energy, execution time, and matched *CapDyn* stage. Tasks and their respective functions are summarized in Table 5.3. ***OP*** stands for "over provisioning'.

Task	Avg Power [mW]	Avg Exec Time [s]	Avg Energy [mJ]	3 stg map D1	3 stg map D2	4 stg map D1	4 stg map D2	5 stg map D1	5 stg map D2	6 stg map D1	6 stg map D2	7 stg map D1	7 stg map D2
$T_{1,E}$	18.00	0.92	16.62	1	1	2	1	2	2	3	3	4	4
$T_{2,E}$	3.40	600.00	2037.54	3	3	4	4	5	5	6	6	7	7
$T_{3,E}$	33.29	13.70	455.91	2	2	3	3	4	4	5	5	6	6
$T_{4,E}$	26.45	0.22	5.87	1	1	1	1	2	2	3	2	4	3
$T_{5,E}$	5.17	7.23	37.40	1	1	2	2	3	3	4	4	5	5
$T_{6,E}$	33.29	68.48	2279.55	3	3	4	4	5	5	6	6	7	7
OP [%]	–	–	–	9.77	12.31	10.53	10.58	10.92	10.12	13.59	12.33	18.60	22.60
$T_{1,L}$	18.00	25.58	460.61	2	2	3	3	4	4	5	5	6	6
$T_{2,L}$	3.40	537.36	1824.81	3	3	4	4	5	5	6	6	7	7
$T_{3,L}$	33.29	41.16	1370.08	3	3	4	4	5	5	6	6	7	7
$T_{4,L}$	26.45	0.22	5.87	1	1	1	1	2	2	3	2	4	3
$T_{5,L}$	5.17	176.96	915.34	3	3	4	4	5	5	6	6	7	7
$T_{6,L}$	33.29	68.48	2279.55	3	3	4	4	5	5	6	6	7	7
OP [%]	–	–	–	42.28	42.98	42.82	42.93	44.88	41.86	48.25	44.92	54.74	58.13

Table 5.5: Per stage capacity/energy vs number of available *CapDYN* stages. Single cap represent the baseline with fixed capacitance/energy.

Stage Number	Stage 1 mF	Stage 1 mJ	Stage 2 mF	Stage 2 mJ	Stage 3 mF	Stage 3 mJ	Stage 4 mF	Stage 4 mJ	Stage 5 mF	Stage 5 mJ	Stage 6 mF	Stage 6 mJ	Stage 7 mF	Stage 7 mJ
3 - D1	35.60	72.09	89.00	504.63	356.00	2299.32	-	-	-	-	-	-	-	-
3 - D2	21.88	106.31	87.50	531.56	350.00	2296.00	-	-	-	-	-	-	-	-
4 - D1	16.18	18.09	35.60	90.18	89.00	522.72	356.00	2317.40	-	-	-	-	-	-
4 - D2	5.44	26.43	21.75	132.13	87.00	554.95	348.00	2309.31	-	-	-	-	-	-
5 - D1	7.83	5.30	16.36	23.59	36.00	96.49	90.00	533.89	360.00	2348.74	-	-	-	-
5 - D2	1.38	6.68	5.50	33.41	22.00	140.33	88.00	568.01	352.00	2342.53	-	-	-	-
6 - D1	3.91	1.68	8.00	7.10	16.73	25.80	36.80	100.32	92.00	547.44	368.00	2402.62	-	-
6 - D2	0.36	1.75	1.44	8.73	5.75	36.68	23.00	148.46	92.00	595.58	368.00	2450.76	-	-
7 - D1	2.02	0.56	4.09	2.31	8.35	7.97	17.45	27.48	38.40	105.24	96.00	571.80	384.00	2507.64
7 - D2	0.09	0.46	0.38	2.28	1.50	9.57	6.00	38.73	24.00	155.37	96.00	621.93	384.00	2557.77
Single Cap	453.60	2286.71	-	-	-	-	-	-	-	-	-	-	-	-

both D1 and D2 topologies and varies the total number of storage stages to quantify the tradeoff between energy granularity and infrastructural power consumption. We restrict the design-space exploration to configurations with 3 to 7 stages, which represent realistic implementation points for the considered capacitor-bank architecture. The application task profiling also includes the scheduler overhead required to communicate with the *CapDYN* energy-state interface. Specifically, each scheduled task requires two read/write operations to the next-task register, and each register operation consumes 1.09 mJ over 0.05 s.

Since both D1 and D2 *CapDYN* banks provide a discrete set of approximately exponentially distributed energy levels across their stages, as shown in Fig. 5.13, we evaluate two task sets with differ-

ent energy distributions. The first task set corresponds to the measured real-world application tasks, whose energy consumptions are also approximately exponentially distributed. The second task set is a synthetic linear task set, obtained by linearly spacing the task energy consumptions between the minimum and maximum energies of the real-world task set. This synthetic workload allows us to evaluate applications whose task-energy distribution does not naturally match the staged energy distribution of the capacitor bank.

Table 5.4 reports the task power, execution time, energy consumption, and task-to-stage mapping obtained according to Sec. 4.4.4. Table 5.5 reports the equivalent capacitance and per-stage energy availability of the evaluated *CapDYN* configurations. For each task-to-stage mapping, we also report the overprovisioning (OP), which quantifies how closely the selected *CapDYN* dimensioning matches the considered task set.

The OP values provide an a priori indication of how well each *CapDYN* configuration fits the considered workload. Since the capacitor bank exposes only a discrete set of available energy levels, energy-admissible execution is expected to be most effective when those levels are close to the profiled task energies. For the measured task set, the approximately exponential task-energy distribution aligns more naturally with the staged energy availability of the D1 and D2 banks, resulting in lower OP values. For the synthetic linear task set, several tasks remain mapped to the highest stages, while the additional low-energy stages are used by only a small fraction of the workload. This leads to larger OP values and suggests that additional stages may provide limited benefit unless their energy levels match dominant task-energy requirements.

To evaluate performance under realistic operating conditions, we use experimentally recorded irradiance traces from a photovoltaic panel dataset [121] as the input power source. The irradiance trace is mapped to the electrical input power of a $200\,\mathrm{cm}^2$ solar panel assuming 20% irradiance-to-electrical conversion efficiency [122]. The resulting power trace is then used as the common underlying input for all CapDYN configurations and reference baselines. The irradiance-to-electrical mapping neglects load-variation-induced effects on the harvested electrical power, which are source- and converter-specific and therefore outside the scope of this work. As the behavior of *CapDYN* is connected to the power input it receives, and not to the upstream conversion dynamics, this fact does not compromise

the aforementioned experiments. To stress the storage dimensioning under a conservative condition, the input power is disabled during task execution. The total test duration is fixed to $T_{\text{test}} = 18000$ s, corresponding to 5 h.

We compare *CapDYN* against four reference baselines, each designed to isolate a specific source of improvement brought by *CapDYN*:

- **S0**: a single monolithic capacitor sized to support the most energy-intensive task, i.e., T_6. This baseline does not implement energy-aware scheduling and includes only the minimum infrastructure required to activate and deactivate the load at the upper and lower voltage thresholds. It therefore exposes the cost of executing tasks without checking whether sufficient energy is available.

- **S1**: a single monolithic capacitor sized to support T_6, augmented with energy-aware scheduling and with the same load activation/deactivation infrastructure as S0. As in *CapDYN*, the load power-gates itself whenever no energy-admissible task can be executed. However, the load is reactivated only when the monolithic capacitor reaches its upper voltage threshold. This baseline isolates the benefit of reconfigurable storage under energy-aware scheduling: unlike S1, *CapDYN* can power the load as soon as the stage required by the next energy-admissible task is available.

- **S2**: a single monolithic capacitor sized to support T_6, with energy-aware scheduling but without dedicated load activation/deactivation circuitry. Instead, an always-on MSP430FR5994 uses its low-power comparator as a software-selectable voltage monitor. The storage capacitor voltage is compared against up to 15 thresholds implemented by a high-value resistive tap network and referenced to a fixed low-power voltage reference. This baseline isolates the tradeoff between dedicated load power-gating hardware and continuously powered control logic, and represents the strongest single-capacitor reference in terms of responsiveness.

- **AB**: an ablation version of *CapDYN* without energy-admissible task dispatch. This baseline uses the same reconfigurable storage hardware as *CapDYN*, but tasks are executed only when

the largest stage is fully charged. It therefore isolates the contribution of the scheduler and energy-state interface from the contribution of the reconfigurable capacitor bank itself.

CapDYN and its ablation version are dimensioned according to the procedure described in Sec. 4.4.4. The single-capacitor baselines are instead sized to support the largest task in the application task set, T_6. As shown in Table 5.5, the single-capacitor baseline requires a larger capacitance than the final stages of *CapDYN*, namely 453.6 mF compared with approximately 348–384 mF. This difference arises because the monolithic capacitor cannot exploit the improved energy extraction enabled by dynamic storage reconfiguration. The single-capacitor baselines therefore exclude the effect of adaptive storage capacity and allow us to separately evaluate load power-gating, threshold-based activation, and energy-aware scheduling.

To interpret the benchmark results, we use a set of complementary metrics that separately quantify completed work, temporal responsiveness, and energy efficiency. This separation is important in battery-free systems because performance is not determined only by the amount of work completed over the full experiment. Since harvested power is intermittent and future energy availability is uncertain, the time at which useful task results become available is also relevant. A system that completes useful work earlier provides higher application value if the input power drops before the end of the trace.

Let x denote the evaluated system, typically *CapDYN*, and let $r \in \{\mathrm{S0}, \mathrm{S1}, \mathrm{S2}, \mathrm{AB}\}$ denote a reference baseline. For each task T_i, $n_{i,x}$ is the number of successful executions under system x. We denote by $t_{i,x}^{\mathrm{end,first}}$ the completion time of the first successful execution of task T_i, and by $t_{i,k,x}^{\mathrm{end}}$ the completion time of the k-th successful execution of task T_i. The total experiment duration is denoted by T_{test}.

We first report the total successful-task throughput N_x,

$$N_x = \sum_{i=1}^{N_T} n_{i,x} \tag{5.21}$$

and the relative throughput gain $G_{x/r}^{\mathrm{thr}}$,

$$G_{x/r}^{\text{thr}} = 100 \cdot \frac{N_x - N_r}{N_r} \tag{5.22}$$

Eq. (5.21) and Eq. (5.22) quantify how many useful task completions are obtained over the same input-power trace. This metric is useful to identify systems that fail to complete tasks or that defer task execution so much that fewer executions occur within the experiment. However, throughput alone does not capture *when* tasks are completed: two systems can achieve the same total number of successful executions while producing useful results at very different times.

To capture when the first useful result of each task becomes available, we report the per-task first-completion shift $\Delta t_{i,x/r}$,

$$\Delta t_{i,x/r} = 100 \cdot \frac{t_{i,x}^{\text{end,first}} - t_{i,r}^{\text{end,first}}}{t_{i,r}^{\text{end,first}}} \tag{5.23}$$

A negative value of Eq. (5.23) means that the first completed result of task T_i is available earlier under *CapDYN* than under the reference. This metric identifies which tasks benefit from exposing intermediate energy states to the scheduler. In particular, it shows whether low- and intermediate-energy tasks can produce useful results before the storage system reaches its maximum-energy state.

We also use an aggregate first-completion advance metric $A_{x/r}$,

$$A_{x/r} = -\frac{1}{|\mathcal{T}_{x,r}|} \sum_{i \in \mathcal{T}_{x,r}} 100 \cdot \frac{t_{i,x}^{\text{end,first}} - t_{i,r}^{\text{end,first}}}{t_{i,r}^{\text{end,first}}} \tag{5.24}$$

where $\mathcal{T}_{x,r}$ is the set of tasks for which both systems have a valid first successful completion. The sign convention in Eq. (5.24) makes the interpretation direct: $A_{x/r} > 0$ means that, on average, *CapDYN* makes the first completed task results available earlier than the reference. While Eq. (5.23) identifies the task-level effects, Eq. (5.24) summarizes the average first-result responsiveness of the entire application.

The first-completion metrics consider only the first successful completion of each task type. They therefore quantify initial responsiveness, but they do not capture repeated executions after the first completion. To measure the cumulative availability of useful work over the entire trace, we introduce the timely successful-task yield Y_x,

$$Y_x = \sum_{i=1}^{N_T} \sum_{k=1}^{n_{i,x}} \left(1 - \frac{t_{i,k,x}^{\text{end}}}{T_{\text{test}}}\right) \tag{5.25}$$

Each successful task contributes a value between 0 and 1. A task completed early contributes with a value that is closer to 1, while a task completed near the end of the trace contributes with a lower value. Failed or missed tasks contribute zero. The corresponding relative timely-yield gain is $G_{x/r}^{Y}$.

$$G_{x/r}^{Y} = 100 \cdot \frac{Y_x - Y_r}{Y_r} \tag{5.26}$$

Timely yield complements both throughput and first-completion advance. Throughput measures how many successful executions are completed, whereas first-completion advance measures when the first useful result of each task type becomes available. Timely yield combines these aspects by accounting for all successful completions and weighting them according to how early they occur. This is particularly relevant for battery-free execution: among systems with similar throughput or similar task-cycle duration, the system with higher timely yield has produced more completed useful work earlier in the trace and is therefore less affected by a hypothetical interruption of the input power happening during the test time. Because Y_x depends on task duration and task distribution, we interpret it within each task set and primarily through the relative gain in Eq. (5.26).

Finally, we distinguish raw efficiency η_x from effective efficiency η_x^{eff}. Raw efficiency is defined as:

$$\eta_x = \frac{E_x^{\text{out}}}{E_x^{\text{in}}} \tag{5.27}$$

where E_x^{in} is the harvested input energy and E_x^{out} is the energy delivered to the load. However, raw efficiency can overestimate useful operation when energy is spent on failed executions. We therefore also report effective efficiency:

$$\eta_x^{\text{eff}} = \frac{E_x^{\text{out}} - E_x^{\text{w}}}{E_x^{\text{in}}} \tag{5.28}$$

where E_x^{w} is the energy wasted in failed task attempts. The distinction between Eq. (5.27) and Eq. (5.28) is especially important

for S0, which does not implement energy-aware scheduling and can therefore start tasks that later fail. For S1, S2, AB, and *CapDYN*, failed executions due to power failures are avoided through energy-aware scheduling, so the comparison focuses mainly on effective-efficiency differences,

$$\Delta\eta_{x/r}^{\text{eff}} = \eta_x^{\text{eff}} - \eta_r^{\text{eff}} \tag{5.29}$$

Eq. (5.29) separates the energy cost of reconfiguration, scheduling, and control infrastructure from the energy wasted by unsuccessful task execution.

5.2.3 Experimental results

The results are organized according to the role of each reference baseline. S0 evaluates the cost of executing tasks without energy-aware scheduling. S1 and S2 quantify the benefit of dynamic storage reconfiguration relative to energy-aware monolithic storages with two different task activation mechanisms. Finally, AB isolates the contribution of energy-admissible task dispatch from the contribution of the reconfigurable storage hardware itself.

Impact of failed executions. The comparison with S0 highlights the cost of executing tasks without first verifying that sufficient energy is available. As shown in Table 5.6, S0 achieves relatively high raw efficiency, equal to 89.09% for the exponential task set and 92.70% for the linear task set. However, its effective efficiency decreases to 63.07% and 77.89%, respectively, because part of the delivered energy is spent on failed task executions. In contrast, *CapDYN* avoids these failed executions and reaches effective efficiencies between 90.15% and 97.17% for the exponential task set, and between 91.19% and 97.46% for the linear task set. This corresponds to effective-efficiency gains of 27.08–34.10% in the exponential case and 13.30–19.57% in the linear case (Tab. 5.9). The absolute LTspice efficiency values should be interpreted as optimistic estimates, since the current capacitor-bank model does not explicitly include capacitor self-discharge or PCB-layout-dependent leakage. However, the relative trends remain meaningful because all *CapDYN* configurations and baselines are evaluated under the same LTspice modeling and power-accounting assumptions.

Table 5.6: Efficiency penalty of the unaware single-capacitor baseline S0 and effective-efficiency comparison with *CapDYN*.

Task set	Bank	N	η_{S0}	$\eta_{\mathrm{S0}}^{\mathrm{eff}}$	$\eta_{CapDYN}^{\mathrm{eff}}$	$\Delta\eta^{\mathrm{eff}}$
Expon.	D1	3	89.09	63.07	95.88	32.82
Expon.	D1	4	89.09	63.07	95.39	32.32
Expon.	D1	5	89.09	63.07	94.53	31.47
Expon.	D1	6	89.09	63.07	93.25	30.18
Expon.	D1	7	89.09	63.07	91.13	28.07
Expon.	D2	3	89.09	63.07	97.09	34.02
Expon.	D2	4	89.09	63.07	97.17	34.10
Expon.	D2	5	89.09	63.07	96.38	33.31
Expon.	D2	6	89.09	63.07	93.57	30.50
Expon.	D2	7	89.09	63.07	90.15	27.09
Linear	D1	3	92.70	77.89	96.56	18.66
Linear	D1	4	92.70	77.89	96.22	18.33
Linear	D1	5	92.70	77.89	95.57	17.67
Linear	D1	6	92.70	77.89	94.65	16.75
Linear	D1	7	92.70	77.89	93.11	15.22
Linear	D2	3	92.70	77.89	97.46	19.56
Linear	D2	4	92.70	77.89	97.41	19.52
Linear	D2	5	92.70	77.89	96.59	18.69
Linear	D2	6	92.70	77.89	94.29	16.40
Linear	D2	7	92.70	77.89	91.19	13.30

Notes. Efficiencies are reported in percent. $\Delta\eta^{\mathrm{eff}} = \eta_{CapDYN}^{\mathrm{eff}} - \eta_{\mathrm{S0}}^{\mathrm{eff}}$ is reported in %.

The same effect seen in effective efficiency is visible in the throughput results of Table 5.7. For the exponential task set, *CapDYN* completes 17 successful tasks, whereas S0 completes 12, corresponding to a throughput gain of 41.67%. For the linear task set, the gain remains positive but is smaller: *CapDYN* completes 13 successful tasks for the D1 configurations and for the 3–5 stage D2 configurations, corresponding to a 30% gain over S0, while the 6–7 stage D2 configurations complete 12 tasks, corresponding to a 20% gain. This reduction is consistent with the poorer match between the linear task-energy distribution and the exponential energy availability of the *CapDYN* banks, as indicated by the larger overprovisioning values in Table 5.4, and with the higher infrastructural consumption of larger D2 configurations. Overall, the comparison with S0 shows that avoiding failed executions improves both energy use and task productivity: a larger fraction of the harvested energy is converted into successful task execution, and more useful work is completed

Table 5.7: Successful-task throughput and relative throughput gain of *CapDYN* with respect to the reference baselines.

Task set	Bank	N_{stg}	N_{CapDYN}	S0	S1	S2	AB
Expon.	D1	3-7	17	41.67	21.43	0.00	21.43
Expon.	D2	3-7	17	41.67	21.43	0.00	21.43
Linear	D1	3-7	13	30.00	8.33	8.33	8.33
Linear	D2	3-5	13	30.00	8.33	8.33	8.33
Linear	D2	6-7	12	20.00	0.00	0.00	0.00

Notes. N_{CapDYN} is the total number of successful task completions under *CapDYN*. The S0, S1, S2, and AB columns report $G^{\text{thr}}_{CapDYN/r}$ in percent.

across the same power-input trace.

Table 5.8 further shows that, with *CapDYN*, the successfully completed tasks are also made available earlier in the trace. Relative to S0, *CapDYN* improves timely successful-task yield by 53.72–59.85% for the exponential task set and by 27.80–35.33% for the linear task set. Thus, the benefit over S0 is not limited to avoiding wasted energy: *CapDYN* also makes successful task results available earlier, increasing the time-weighted amount of useful work completed over the trace.

Impact of dynamic storage reconfiguration. The comparisons with S1 and S2 isolate the effect of replacing energy-aware monolithic storage with the reconfigurable *CapDYN* bank. Both baselines use energy-aware scheduling, so failed task attempts due to insufficient stored energy are avoided by design. However, they differ in how the load resumes execution. S1 uses the same load activation/deactivation infrastructure as S0 and therefore reactivates the load only when the monolithic capacitor reaches its upper voltage threshold. S2 instead keeps an always-on MCU active and uses its programmable voltage monitor to resume execution once sufficient energy has been stored. Thus, S1 directly tests the reactivation-latency benefit of *CapDYN*'s staged storage: both S1 and *CapDYN* power-gate the load, but S1 must wait for the full monolithic capacitor to reach its upper threshold before execution can resume, whereas *CapDYN* can reactivate the load at an intermediate energy state. S2 provides the strongest single-capacitor responsiveness baseline because its always-on MCU can resume execution without waiting for full-capacitor recharge.

Table 5.8: Timely successful-task yield gain of *CapDYN* relative to the reference baselines. Positive values indicate higher time-weighted useful task yield.

Task set	Bank	N	Y_{CapDYN}	S0	S1	S2	AB
Expon.	D1	3	8.175	58.07	26.47	3.69	21.33
Expon.	D1	4	8.198	58.52	26.83	3.99	21.77
Expon.	D1	5	8.224	59.01	27.23	4.31	22.57
Expon.	D1	6	8.197	58.50	26.82	3.97	22.39
Expon.	D1	7	8.176	58.09	26.49	3.71	23.32
Expon.	D2	3	8.240	59.33	27.48	4.52	37.18
Expon.	D2	4	8.267	59.85	27.90	4.86	22.77
Expon.	D2	5	8.240	59.32	27.48	4.52	22.78
Expon.	D2	6	8.220	58.94	27.17	4.26	23.68
Expon.	D2	7	7.950	53.72	22.99	0.84	23.00
Linear	D1	3	5.543	35.11	19.75	-0.93	7.49
Linear	D1	4	5.534	34.88	19.55	-1.10	7.56
Linear	D1	5	5.524	34.64	19.34	-1.27	7.53
Linear	D1	6	5.504	34.15	18.91	-1.63	7.70
Linear	D1	7	5.466	33.24	18.10	-2.30	8.27
Linear	D2	3	5.552	35.33	19.96	-0.76	7.41
Linear	D2	4	5.538	34.98	19.64	-1.02	7.40
Linear	D2	5	5.512	34.34	19.08	-1.49	7.72
Linear	D2	6	5.332	29.96	15.19	-4.71	7.81
Linear	D2	7	5.243	27.80	13.28	-6.29	8.38

Notes. Y_{CapDYN} is the absolute timely successful-task yield of *CapDYN*. The S0, S1, S2, and AB columns report $G^{Y}_{CapDYN/r}$ in percent.

As shown in Table 5.9, *CapDYN* improves effective efficiency relative to both energy-aware single-capacitor baselines. Against S1, the effective-efficiency gain is 4.52–11.54% for the exponential task set and 2.16–8.43% for the linear task set. Against S2, the corresponding gains are 3.78–10.80% and 1.36–7.62%. These results show that the benefit of *CapDYN* is not only due to energy-aware scheduling: even when failed executions are already avoided, exposing multiple usable storage states improves how harvested energy is converted into completed useful work. The gain decreases for larger configurations, especially for D2, because additional stages increase infrastructural consumption.

The timing results reflect the different reactivation mechanisms of the two monolithic baselines. Relative to S1, *CapDYN* consistently benefits from exposing intermediate energy states, because S1 must wait for the monolithic capacitor to reach its activation thresh-

Table 5.9: Effective-efficiency comparison of *CapDYN* against the energy-aware baselines and ablation.

Task set	Bank	N	$\eta^{\text{eff}}_{CapDYN}$	ΔS0	ΔS1	ΔS2	ΔAB
Expon.	D1	3	95.88	32.81	10.25	9.51	-0.08
Expon.	D1	4	95.39	32.32	9.75	9.02	-0.07
Expon.	D1	5	94.53	31.46	8.90	8.17	-0.01
Expon.	D1	6	93.25	30.18	7.61	6.88	-0.05
Expon.	D1	7	91.13	28.06	5.50	4.76	-0.06
Expon.	D2	3	97.09	34.02	11.46	10.72	-0.06
Expon.	D2	4	97.17	34.10	11.54	10.80	-0.12
Expon.	D2	5	96.38	33.31	10.74	10.01	-0.22
Expon.	D2	6	93.57	30.50	7.94	7.20	-0.15
Expon.	D2	7	90.15	27.08	4.52	3.78	-0.22
Linear	D1	3	96.56	18.67	7.52	6.72	-0.10
Linear	D1	4	96.22	18.33	7.19	6.39	-0.03
Linear	D1	5	95.57	17.68	6.54	5.73	-0.03
Linear	D1	6	94.65	16.76	5.62	4.81	-0.03
Linear	D1	7	93.11	15.22	4.08	3.28	-0.01
Linear	D2	3	97.46	19.57	8.43	7.62	-0.01
Linear	D2	4	97.41	19.52	8.38	7.58	-0.02
Linear	D2	5	96.59	18.70	7.56	6.75	-0.03
Linear	D2	6	94.29	16.40	5.26	4.45	-0.05
Linear	D2	7	91.19	13.30	2.16	1.36	-0.12

Notes. $\eta^{\text{eff}}_{CapDYN}$ is reported in percent. ΔS0, ΔS1, ΔS2, and ΔAB report effective-efficiency differences in % relative to the corresponding reference.

old before the load can resume. Table 5.10 shows that *CapDYN* achieves an aggregate first-completion advance of 35.57–41.03% for the exponential task set and 20.23–24.55% for the linear task set. The corresponding timely-yield gains are 22.99–27.90% and 13.28–19.96%, respectively, as reported in Table 5.8. The comparison with S2 is stricter: for the exponential task set, *CapDYN* still achieves an aggregate first-completion advance of 18.66–28.74%, with timely-yield gains of 0.84–4.86%. This shows that, when task energies match the staged energy distribution, *CapDYN* can outperform even an always-on energy-aware monolithic baseline S2.

The linear task set exposes the workload-matching limitation of dynamic storage reconfiguration. Relative to S2, the aggregate first-completion advance becomes small and can become negative, reaching −3.94% for the 7-stage D2 bank, while the timely-yield gain ranges from −6.29% to −0.76%. This occurs because the linearly spaced task energies do not align well with the approximately

Table 5.10: Aggregate first-completion advance of *CapDYN* relative to the reference baselines. Positive values indicate earlier first successful task completion under *CapDYN*.

Task set	Bank	N	S0	S1	S2	AB
Expon.	D1	3	40.04	36.14	18.66	30.47
Expon.	D1	4	40.92	36.98	19.78	31.43
Expon.	D1	5	41.60	37.66	21.83	32.62
Expon.	D1	6	40.69	36.74	20.49	32.07
Expon.	D1	7	39.59	35.57	19.05	32.55
Expon.	D2	3	42.74	38.78	22.84	33.10
Expon.	D2	4	45.00	41.03	28.74	35.89
Expon.	D2	5	44.91	40.94	28.61	35.91
Expon.	D2	6	44.69	40.70	28.29	36.20
Expon.	D2	7	41.00	36.93	23.18	34.39
Linear	D1	3	35.22	23.80	0.76	12.61
Linear	D1	4	34.91	23.48	0.26	12.55
Linear	D1	5	34.44	23.03	-0.40	12.33
Linear	D1	6	33.67	22.26	-1.47	12.24
Linear	D1	7	32.51	21.11	-3.01	12.73
Linear	D2	3	35.97	24.55	1.84	12.90
Linear	D2	4	35.97	24.54	1.83	13.04
Linear	D2	5	35.58	24.13	1.27	13.11
Linear	D2	6	33.59	21.94	-1.61	13.13
Linear	D2	7	31.91	20.23	-3.94	13.29

Notes. Values report $A_{CapDYN/r}$ in percent. Positive values indicate earlier first task completion under *CapDYN*. AB denotes the ablation without energy-admissible task dispatch.

exponential stage-energy distributions exposed by D1 and D2. As shown in Tab. 5.4, OP increases from 9.77–22.60% for the exponential task set to 42.28–58.13% for the linear task set. Most linear tasks are mapped to the largest or second-largest stages: in the 7-stage banks, $T_{2,L}$, $T_{3,L}$, $T_{5,L}$, and $T_{6,L}$ are mapped to stage 7, $T_{1,L}$ is mapped to stage 6, and only $T_{4,L}$ uses a lower stage. Consequently, for the high-energy linear tasks, *CapDYN* behaves more similarly to the ablation baseline, because both systems must wait for the bank to reach a high-energy state before these tasks can execute. This is reflected in Table 5.11, where the first-completion shifts relative to AB are close to zero for most high-energy linear tasks, while the remaining temporal gains are concentrated on T_1 and T_4.

Together, these results show that the system-level benefit of *CapDYN* depends on the match between workload energy demand and available stage-energy levels. Energy-admissible execution is most

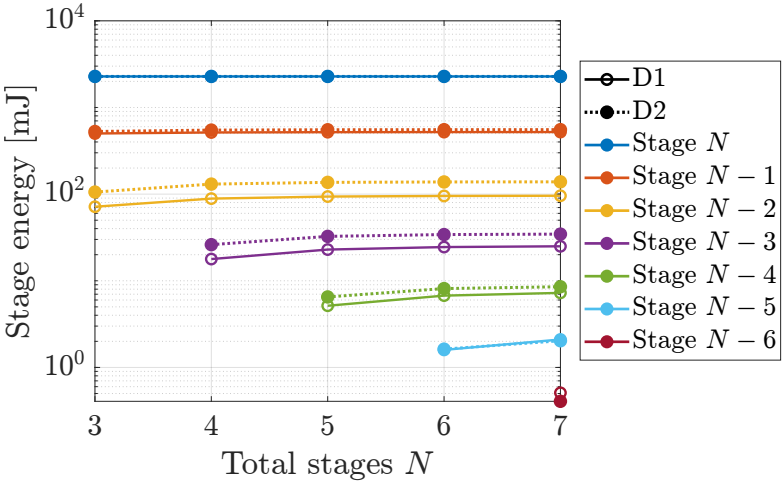


Figure 5.13: Stage-energy availability of *CapDYN* as a function of capacitor-bank topology and total number of stages N. Solid lines and empty markers denote the D1 bank, while dotted lines and full markers denote the D2 bank; colors identify the relative stage position from stage N down to stage $N-6$.

effective when the workload can exploit intermediate storage states, as in the exponential task set. When most tasks remain close to the highest stages, as in the linear task set, additional low-energy stages are underutilized and the temporal advantage over S2 decreases.

This workload dependence also clarifies why reducing OP is not simply a matter of increasing the number of stages. Adding stages gives the scheduler more possible execution points, but OP decreases only when the added or shifted stages move closer to the task-energy requirements that dominate the workload. In the evaluated banks, increasing the number of stages mainly improves granularity at the lower end of the available-energy range, while the gaps between the highest-energy stages remain large, as shown in Fig. 5.13. Thus, additional low-energy stages provide limited benefit when most tasks remain close to the highest-energy stages.

This is the case for the linear task set. Most linear tasks require 460.61–2279.55 mJ. Across the evaluated D1 and D2 banks, the largest stage provides 2296.00–2557.77 mJ and the second-largest stage provides 504.63–621.93 mJ, while the third-largest stage is

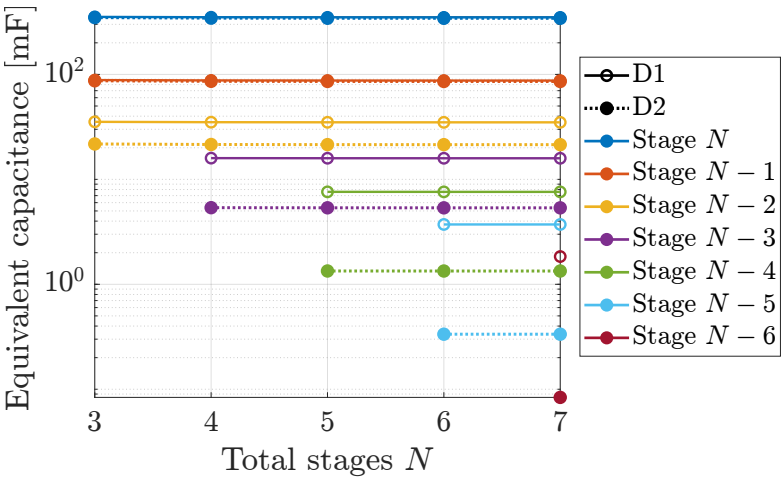


Figure 5.14: Stage capacitance of *CapDYN* as a function of capacitor-bank topology and total number of stages N. Solid lines and empty markers denote the D1 bank, while dotted lines and full markers denote the D2 bank; colors identify the relative stage position from stage N down to stage $N-6$.

only 72.09–155.37 mJ, as reported in Table 5.5. Therefore, stages below the second-largest level are too small for most linear tasks and can only be exploited by the low-energy task $T_{4,L}$, which requires 5.87 mJ. This explains why increasing the number of stages does not necessarily improve temporal performance and why S2 can remain competitive in some linear-task-set *CapDYN* configurations.

Impact of energy-admissible dispatch. The comparison with AB isolates the contribution of energy-admissible task dispatch. Since AB uses the same reconfigurable storage hardware as *CapDYN*, efficiency differences between the two systems are expected to be small. Table 5.9 confirms this behavior: in the available configurations, the effective-efficiency difference relative to AB remains between approximately -0.01 and -0.22 %. Thus, enabling energy-admissible dispatch does not substantially change the energy efficiency of the underlying reconfigurable storage hardware.

The benefit of dispatch is instead primarily temporal. Table 5.7 shows that, for the exponential task set, *CapDYN* completes 17

tasks while AB completes 14, corresponding to a throughput gain of 21.43%. For the linear task set, the gain is smaller, typically 8.33%, and becomes 0% for the largest available D2 configurations, namely the 6- and 7-stage banks. This reduction is due to both workload mapping and bank overhead: as previously mentioned, most linear tasks are mapped to high stages, which limits the benefit of intermediate-stage dispatch, while the increased infrastructural power consumption of larger D2 banks further offsets the advantage of finer energy granularity.

The aggregate first-completion results in Table 5.10 show that dispatch significantly improves responsiveness relative to AB. Relative to AB, *CapDYN* achieves an aggregate first-completion advance of 30.47–36.20% for the exponential task set and 12.24–13.29% for the linear task set. Similarly, Table 5.8 shows positive timely-yield gains of 21.33–37.18% and 7.40–8.38%, respectively. The smaller gains for the linear task set are consistent with the workload-matching limitation discussed above: fewer tasks can benefit from intermediate-stage dispatch. Even so, the gains remain positive, demonstrating that energy-admissible dispatch increases the amount of useful work available earlier in the trace, even when the underlying storage hardware is unchanged.

The per-task comparison in Table 5.11 identifies which tasks drive this temporal benefit. For the exponential task set, *CapDYN* substantially advances low- and intermediate-energy tasks. For example, the first completion of T_1 occurs 76.16–96.33% earlier than in AB, while T_3, T_4, and T_5 also show large negative shifts. By contrast, T_2 and T_6, which require the largest energy reserves, show shifts close to zero because both *CapDYN* and AB must wait for high-energy storage states before these tasks can complete. This confirms the intended function of energy-admissible dispatch: it does not accelerate tasks that inherently require the largest storage state, but it avoids delaying lower-energy tasks until that state is reached.

For the linear task set, the per-task benefit is more concentrated. The first completion of T_1 is advanced by approximately 54.53–57.23%, and T_4 by approximately 19.36–23.25%. The remaining tasks show shifts close to zero, and in a few cases slightly positive values, because most linear-task energies are mapped to the highest stages. Even in this mismatched workload, however, *CapDYN* improves the timing of the tasks that can exploit intermediate stages,

which explains the positive aggregate first-completion and timely-yield gains relative to AB.

Best configurations and design implications. For the evaluated traces and task sets, the D2 bank gives the best overall configurations, but with different stage counts, as illustrated in Figs. 5.15 and 5.16. For the exponential task set, D2 with $N = 4$ provides the best overall tradeoff, reaching the highest effective efficiency, 97.17%, and the highest timely successful-task yield, $Y_{CapDYN} = 8.267$. It also gives the largest aggregate first-completion advance against the single-capacitor baselines, equal to 45.00%, 41.03%, and 28.74% relative to S0, S1, and S2, respectively. For the linear task set, D2 with $N = 3$ is preferable, reaching the highest effective efficiency, 97.46%, and the highest timely successful-task yield, $Y_{CapDYN} = 5.552$. It also provides the largest aggregate first-completion advance against the single-capacitor baselines, equal to 35.97%, 24.55%, and 1.84% relative to S0, S1, and S2, respectively.

Although larger D2 banks can slightly improve some AB-only timing metrics, they are not the best overall choice for the two considered task sets. In the LTspice model, D2 infrastructural consumption increases rapidly with stage count because the number of switches grows with a base-2 dependence, whereas D1 scales approximately linearly, as shown in Tab. 5.12[1] Therefore, when the workload does not effectively exploit the additional D2 stages, the extra overhead can outweigh the benefit of finer energy granularity. For applications whose task-energy distribution can benefit from a larger number of useful stages, D1 could become preferable because its infrastructural consumption scales more gently with stage count. This reinforces the need for workload-aware dimensioning, as the optimal topology is determined by the joint effect of task-to-stage matching, stage count, and infrastructural power consumption.

Summary. Overall, the LTspice results show that *CapDYN* provides two complementary application-level benefits. Relative to the single-capacitor baselines, it improves effective efficiency, throughput, and responsiveness by avoiding failed executions and exposing

[1]The LTspice values are conservative with respect to the hardware measurements (Fig. 5.6), as they slightly overestimate infrastructural power, while preserving the same topology- and stage-dependent scaling trends

intermediate storage states to the runtime. Relative to AB, the results show that the reconfigurable bank alone is not sufficient: the energy-state interface must be coupled with energy-admissible dispatch to convert intermediate storage states into earlier useful task completions.

These results also connect the LTspice evaluation with the hardware measurements. At circuit level, *CapDYN* reduces leftover energy and exposes staged energy availability, while introducing infrastructural and bank-side losses. At application level, the resulting tradeoff depends on how effectively the workload exploits the exposed stages. The exponential task set therefore achieves the strongest gains, whereas the linear task set shows that, when most task energies remain close to the highest stages, additional low-energy stages are underutilized and the advantage over the always-on monolithic baseline S2 decreases.

The analytical task-to-stage mapping and the hardware efficiency metrics describe complementary properties. Overprovisioning measures the nominal energy assigned in excess of a task requirement, which may remain available for later execution and is therefore not necessarily dissipated. Infrastructural, bank-side, and equalization losses instead represent the energy costs of physically producing and operating the selected storage configuration. Application-specific dimensioning should consequently evaluate improvements in task-to-stage matching against the topology-, stage-, and current-dependent hardware costs characterized in this work. The formal integration of these criteria into a unified selection procedure is left for future work.

5.3 Discussion

In this section, we discuss the limits of *CapDYN* and future work.
Hardware implementation: The relationship between capacitors to make the D1 bank design more efficient is to be further studied. The D1 bank is a promising solution, as D1 is easier to scale up in stage number with respect to D2, and D1 can be adapted to loads with tighter voltage ranges. The development of a new capacitor dimensioning solution could fully remove D1's small equalization transients.

Besides optimizing the capacitor bank, we can improve over-

Table 5.11: Per-task first-completion shift of *CapDYN* relative to the ablation baseline AB. Negative values indicate earlier first task completion under *CapDYN*.

Task set	Bank	N	T_1	T_2	T_3	T_4	T_5	T_6
Expon.	D1	3	-76.16	-0.04	-37.18	-34.71	-34.68	-0.03
Expon.	D1	4	-75.78	-0.01	-37.01	-41.78	-33.97	-0.03
Expon.	D1	5	-82.35	-0.02	-37.08	-41.77	-34.05	-0.41
Expon.	D1	6	-81.66	0.01	-36.39	-40.87	-33.46	-0.02
Expon.	D1	7	-82.01	-0.02	-35.54	-38.73	-38.72	-0.28
Expon.	D2	3	-79.34	-0.25	-37.00	-41.02	-40.85	-0.14
Expon.	D2	4	-96.33	-0.11	-36.68	-41.64	-40.44	-0.13
Expon.	D2	5	-96.11	-0.05	-36.92	-41.76	-40.58	-0.03
Expon.	D2	6	-95.68	-0.07	-37.63	-42.19	-41.28	-0.34
Expon.	D2	7	-94.06	0.03	-35.72	-39.08	-37.21	-0.31
Linear	D1	3	-55.54	-0.39	-0.17	-19.41	-0.13	-0.03
Linear	D1	4	-54.91	-0.40	-0.14	-19.47	-0.35	-0.03
Linear	D1	5	-54.53	-0.08	-0.02	-19.36	-0.01	-0.00
Linear	D1	6	-54.63	0.53	0.07	-19.65	0.21	0.03
Linear	D1	7	-56.29	0.42	0.06	-20.71	0.08	0.06
Linear	D2	3	-56.27	-1.08	-0.26	-19.43	-0.32	-0.02
Linear	D2	4	-56.31	-1.37	-0.40	-19.84	-0.31	-0.02
Linear	D2	5	-55.97	-0.80	-0.57	-20.34	-0.88	-0.08
Linear	D2	6	-55.96	-0.12	-0.01	-22.67	0.01	-0.00
Linear	D2	7	-57.23	0.46	0.13	-23.25	0.12	0.05

Notes. Values report the first-completion shift $\Delta t_{i,CapDYN/AB}$ in percent. AB denotes the ablation without energy-admissible task dispatch.

all energy efficiency and cold start times by reducing infrastructural power consumption, thus shrinking the required size of the backup capacitor. To begin with, charge-pump like voltage boosting could be more efficient than including a *harvester-to-DC* conditioning stage, in the case of low-voltage DC sources. On the other hand, charge-pump boosting is only efficient if no voltage equalization transients appear, and harvester-to-DC feature integrated maximum power point tracking algorithms, which can make up for the additional power expenditure by harvesting more energy from the source. Furthermore, a VLSI implementation has the potential to significantly reduce power requirements and the overall area occupied by the electronic circuitry. This advancement would expand the range of applications to even lower power and more compact scenarios, further reducing *CapDYN*'s power consumption. This would expand the lowest current value we can harvest, without sacrificing reactivity. For the sake of modularity and design freedom,

Table 5.12: LTspice infrastructural power consumption of the single-capacitor baselines and of the reconfigurable *CapDYN* hardware used by both *CapDYN* and AB. The values are conservative with respect to the hardware measurements of Sec. 5.1.4, as they slightly overestimate infrastructural power while preserving the same topology- and stage-dependent scaling trends.

System	N	P_{INFR} **(μW)**
S0	–	4.155
S1	–	4.186
S2	–	4.100
D1	3	11.925
D1	4	12.309
D1	5	12.659
D1	6	13.029
D1	7	13.367
D2	3	12.463
D2	4	13.099
D2	5	14.755
D2	6	17.980
D2	7	29.206

we could also integrate all the components needed to expand one stage into one single chip.

Resilience to prolonged power failures could also be improved through a different design of the capacitor-bank switches. As discussed in Sec. 4.4.1, in the present implementation the bank configuration can only be guaranteed while the analog-switch supply remains within its operating range. Once this supply falls below the switches' minimum operating voltage, their channel state can no longer be guaranteed and the energy remaining in the main bank may eventually be lost. A possible direction is to adopt state-retaining switch techniques such as those employed by Capybara [21], where a small latch capacitor locally stores the switch-control state while the control unit is unpowered. Applied to *CapDYN*, such an approach could decouple preservation of the bank configuration from continuous powering of the control infrastructure, allowing the energy stored in the main bank to remain available throughout substantially longer outages and to be recovered when harvested power becomes available again. However, unlike Capybara's independently switched capacitor banks, *CapDYN*'s switches jointly define coordinated series–parallel configurations. A state-retaining implementation would therefore need to preserve a valid

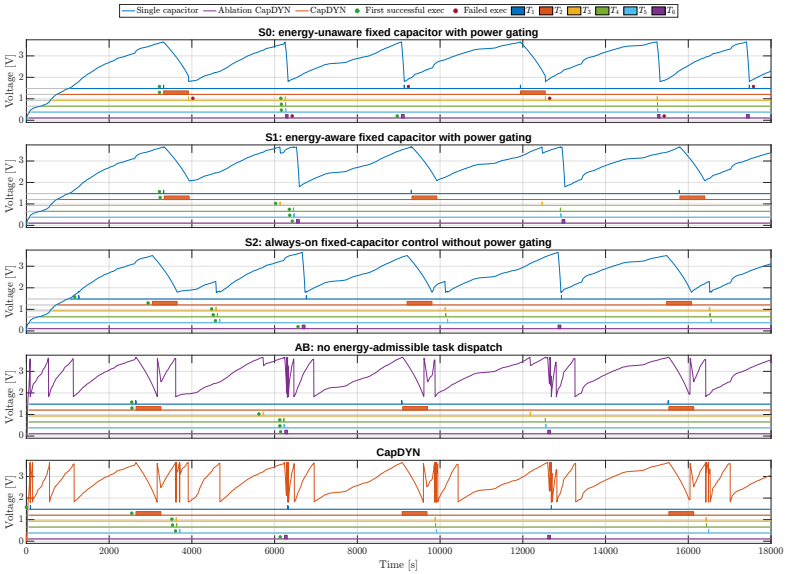


Figure 5.15: Voltage and task-execution comparison for the best-performing *CapDYN* configuration on the exponential task set. The figure compares the single-capacitor baselines S0, S1, and S2, the ablation baseline AB, and *CapDYN* with a four-stage D2 bank. This configuration provides the best overall tradeoff for the exponential task set, combining the highest effective efficiency and timely successful-task yield with earlier task completions relative to the single-capacitor baselines and to the ablation.

bank-stage configuration and resynchronize the retained physical stage with the scaling logic upon power recovery. The resulting switch cells would also need to satisfy the ultra-low leakage, low on-resistance, bidirectional conduction, and operating-voltage requirements of the present design. A large number of stages could also have interesting applications, as it could provide both fine-grained capacitance control and pre-regulation. This would allow designers to reduce charging overheads by perfectly matching requested task energy to the energy available at a certain stage.

Adaptation to unknown or dynamic workloads: In the present implementation, *CapDYN* assumes that task energy requirements

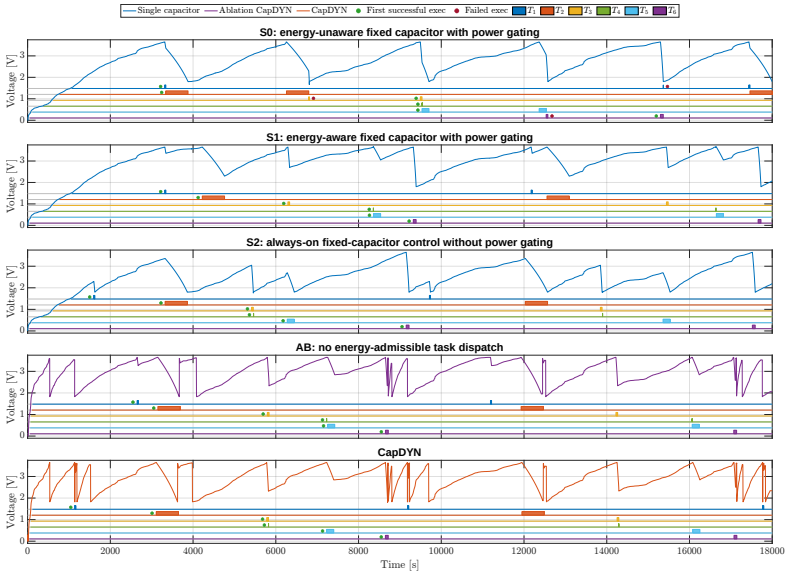


Figure 5.16: Voltage and task-execution comparison for the best-performing *CapDYN* configuration on the linear task set. The figure compares the single-capacitor baselines S0, S1, and S2, the ablation baseline AB, and *CapDYN* with a three-stage D2 bank. This configuration provides the best overall tradeoff for the linear task set, combining the highest effective efficiency and timely successful-task yield with earlier task completions relative to the single-capacitor baselines and to the ablation.

are known a priori through offline profiling, and that the total number of physical stages is fixed at design time. A future implementation could extend *CapDYN* to support initially unknown or dynamically changing workloads by combining online task-energy characterization with a programmable *effective* number of stages.

For an initially unknown workload, the first application cycles could be used as an online training phase. Tasks could initially be assigned conservatively to the largest available energy state, while their energy requirements are characterized either through on-board energy estimation or through trial-and-error execution outcomes. If the available configuration proved insufficient to sustain a task, additional stages could be enabled, up to the maximum physically

implemented bank size. Once task-energy estimates are available, an optimization procedure analogous to the task-to-stage mapping of Sec. 4.4.4 could jointly determine the task-to-stage assignment and the effective number of enabled stages, with the objective of reducing overprovisioning and maximizing end-to-end efficiency while accounting for the infrastructural cost of the enabled stages.

For workloads whose energy requirements change over time, this training process could be repeated when significant deviations from the previously learned task-energy behavior are detected. The resulting measurements could progressively build an energy-consumption model that relates each task's requirements to relevant influence factors, such as task size, temperature, component aging, or previous execution outcomes. This model could then be used to predict both the appropriate *CapDYN* stage for each task and the effective number of stages required by the current workload. If the evolution of these factors can be predicted with sufficient accuracy, explicit retraining could eventually be reduced or avoided by directly forecasting the required task-to-stage mapping and storage configuration.

The achievable online adaptation would nevertheless remain bounded by the physical hardware. In particular, the capacitor values mounted on the board and the maximum implemented stage count constrain the set of energy levels that can be generated, and therefore limit how closely an unknown or changing task-energy distribution can be matched after deployment. Making the effective stage count programmable would also require additional switching hardware, increasing circuit complexity and infrastructural power consumption.

Energy and Task-based Intermittent Computing: *CapDYN* employs the task-based intermittent programming model, which fits best with the staged capacitor bank. The proper mapping of the tasks' energy consumption to the stages' energy availability, fully eliminates the power-failure interrupts during task execution. Power outages can only happen between tasks while waiting for charging, which introduces neither a non-termination issue nor memory inconsistency. Moreover, depending on the application, *CapDYN*'s energy state interface allows for scheduling the flow and order of the tasks according to the available energy in the current stage of the capacitor bank. For instance, a task with a high priority can be executed earlier than other minor tasks. In this sense, also time-constrained tasks can take a chance to run intermittently. Besides,

we could also develop a scheduler that prioritizes certain tasks depending on the power input level by measuring the charging times at the load side through persistent timekeepers [100], or within *CapDYN* itself. We leave the implementation of such a scheduler for future work.

***CapDYN* with Checkpoints:** Although the task-based model is more directly aligned with *CapDYN* due to its energy-state interface, checkpointing approaches can also be combined with the proposed storage support. In this case, checkpoint intervals, task fragments, or preemptive segments should be treated as energy-profiled execution units, and the energy required to save and restore execution state must be included explicitly in the admissibility condition. For static checkpointing, the programmer must consider the first-stage energy availability of *CapDYN* to determine the placement of checkpoints. Dynamic checkpointing can instead use the stage information coming from *CapDYN* hardware: if the current stage has enough energy to complete the next checkpoint interval and the associated save/restore overhead, the checkpoint operation can be omitted or delayed at runtime to reduce energy consumption.

5.4 Conclusions

This chapter experimentally evaluated the *CapDYN* architecture developed in Chapter 4 at both circuit and system level. The proposed architecture replaces the conventional monolithic storage capacitor with a dynamic capacitor bank whose equivalent capacitance is autonomously reconfigured during charge and discharge. This design allows the storage subsystem to operate not only as an energy buffer which minimizes leftover energy, but also as a discrete energy-state interface through which the runtime can enforce energy-admissible task execution. The evaluation considered both the physical behavior of the reconfigurable capacitor bank and the application-level benefits obtained when its energy-state interface is coupled to the energy-admissible runtime.

The hardware evaluation validated the autonomous switching mechanism on D1 and D2 capacitor-bank prototypes. The measurements showed that *CapDYN* progresses correctly through its charging and discharging stages, keeps the load within the target voltage range, and reduces leftover energy from 23.67% of the input

energy in the fixed-capacitor baseline to 0.16%–5.57% across the measured *CapDYN* configurations. At the same time, the measurements quantified the cost of reconfiguration: the infrastructure remains below 10 μW, ranging from 4.15 μW for the 3-stage D1 configuration to 9.07 μW for the 5-stage D2 configuration, while the bank introduces aggregate parasitic losses consistent with contributions from switch-path resistance, capacitor/switch leakage, and switching equalization transients. These losses affect the two topologies differently: D1 exhibits slight equalization transients when capacitor groups are parallelized during charge-stage increases, whereas D2 nominally avoids these transients but uses more switches, increasing conduction paths and leakage opportunities as the number of stages grows. Thus, *CapDYN* remains an ultra-low-power architecture even in a discrete-component implementation, but it should not be interpreted simply as a more efficient capacitor in isolation. Its main benefit emerges at system level, where the application can use the exposed energy-state information to decide when tasks should run and exploit intermediate storage states for energy-admissible execution.

The LTspice application-level evaluation showed how this circuit-level characterization translates into system-level benefits in a complete battery-free sensing scenario, where sensing, communication, imaging, and impedance-measurement tasks are scheduled under realistic harvested-energy traces. Relative to the single-capacitor baselines, *CapDYN* improves effective efficiency by 27.09–34.10% over S0 for the better-matched task set and by 13.30–19.56% for the less-matched task set. Even against the energy-aware monolithic baselines, where failed executions are already avoided, *CapDYN* improves effective efficiency by up to 11.54% over S1 and 10.80% over S2. The same trend appears in responsiveness: aggregate first-completion advance reaches 45.00%, 41.03%, and 28.74% relative to S0, S1, and S2, respectively, for the better-matched task set, while the corresponding maxima for the less-matched task set are 35.97%, 24.55%, and 1.84%. The comparison with AB shows that these gains do not come from the reconfigurable bank alone. Since AB uses the same storage hardware without energy-admissible dispatch, the effective-efficiency difference remains negligible, between -0.01% and -0.22%. The benefit is instead temporal: *CapDYN* advances aggregate first completions by 30.47–36.20% over AB for the better-matched task set and by 12.24–13.29% for the less-matched

task set. Thus, the LTspice evaluation shows that the hardware and runtime are both necessary: the bank exposes intermediate energy states, and the scheduler turns them into earlier useful task completions.

The combined hardware and simulation results also demonstrate that *CapDYN* dimensioning is inherently workload dependent. Increasing the number of stages improves energy granularity only when the additional energy levels align with the task-energy distribution. Otherwise, the added stages may remain underutilized while still increasing infrastructural overhead, particularly for D2, whose implementation cost grows rapidly with stage count. For the evaluated traces and task sets, compact D2 configurations provide the best overall tradeoff: $N = 4$ is preferable for the better-matched task set, whereas $N = 3$ is preferable when most task energies remain concentrated near the highest storage levels. This result indicates that no single topology or stage count is optimal in general. The appropriate configuration must instead be selected by jointly considering task-energy clusters, capacitor-bank topology, stage count, and infrastructural power consumption.

Overall, this chapter shows that the energy storage for battery-free systems can be formulated as a hardware–software co-design problem. The hardware provides autonomous dynamic capacitance and exposes discrete energy states, while the runtime uses those states to admit tasks only when sufficient energy is available. The system-level benefit emerges from the interaction between these two components: the capacitor bank creates intermediate energy states, and the scheduler converts them into earlier, failure-free task execution.

Future work will proceed in three directions. First, the LTspice model will be refined by adding capacitor-leakage terms and PCB-layout-dependent parasitic effects, improving the prediction of absolute efficiency while preserving its role as a design-space exploration tool. Second, workload-aware sizing algorithms will be developed to place energy stages near dominant task-energy clusters while accounting for capacitor availability and implementation overhead. Third, integrated implementations will be investigated to further reduce power consumption, footprint, and parasitic losses, extending autonomous dynamic capacitance to smaller and lower-power battery-free platforms.

Chapter 6

Conclusions

This thesis investigated energy-management circuits and architectures for battery-free embedded systems powered by sustainable ambient energy sources, with particular attention to Plant Microbial Fuel Cells (PMFCs). The starting point of the work was the mismatch between the way energy is produced by such sources and the way it is consumed by embedded workloads. PMFCs provide low and slowly varying power, affected by biological, electrochemical, and environmental conditions, whereas sensing, computation, communication, and monitoring tasks require reliable bursts of usable energy. Addressing this mismatch requires looking beyond average harvested power alone: it requires understanding how energy is extracted from the source, how it is buffered, and when it becomes available to the runtime for useful task execution.

The thesis addressed this problem along the energy path of a battery-free node. On the source side, it studied controlled intermittent extraction from PMFCs, with the goal of improving the amount of energy that can be harvested from a sustainable but difficult-to-control source. On the storage side, it developed *CapDYN*, an autonomous dynamic-capacitance architecture that adapts the energy buffer to the needs of intermittent execution. These two contributions operate at different points of the system, but they share the same objective: increasing the fraction of scarce harvested energy that becomes completed battery-free tasks.

The first contribution was the design and evaluation of a two-electrode intermittent-harvesting interface for PMFCs. The pro-

posed interface imposes extraction and recovery intervals explicitly, instead of allowing them to emerge from capacitor thresholds, converter voltage windows, or load demand. In this way, harvesting frequency and duty cycle become controlled circuit-level variables, while the PMFC remains connected to a converter-based energy-delivery path. This bridges part of the gap between electrochemical studies, which can precisely control loading protocols but are not directly deployable as harvesters, and embedded harvesting systems, which are deployable but often do not independently control the intermittent timing.

The experimental results show that controlled intermittent extraction is a promising strategy for PMFC-powered systems. Preliminary timing tests indicated that seconds-scale periods, especially around 5–10 s in the tested cells, were favorable. In the longer 9-day experiment, reducing the intermittent period from 20 s to 10 s coincided with average-power increases of 159% and 248% in the two intermittently loaded PMFCs, while the continuously loaded reference increased by 31% over the same interval. Relative to the continuous reference, the two intermittent cells reached +185% and +57% during the final 10 s phase. These results support the hypothesis that suitable recovery and extraction timing can improve PMFC energy harvesting. At the same time, the increase observed in the continuous reference confirms that biological and environmental dynamics remain important confounding factors. Therefore, given also the limited number of cells, the results should be interpreted as evidence of the potential of controlled intermittency, rather than as a complete isolation of the biological mechanisms responsible for the observed gains.

The PMFC investigation also showed that power extraction and state observation should be treated as coupled aspects of PMFC energy management. Electrochemical Impedance Spectroscopy (EIS) was used to investigate whether different loading histories leave measurable signatures in the PMFC response. The impedance results showed loading-history-dependent differences in the measured anode and cathode responses, while hydration can also strongly influence the spectra. This supports the use of impedance-based observation as a non-destructive diagnostic tool for PMFC-powered nodes, but also highlights the need for careful interpretation: the measured impedance reflects coupled effects of electrode state, hydration, microbial activity, and loading history. Future PMFC power managers

should therefore combine power-based measurements with electrochemical observables, so that operating voltage, intermittent timing, and diagnostic measurements can be jointly optimized.

The second contribution of the thesis was *CapDYN*, an autonomous dynamic-capacitance architecture for battery-free intermittent systems. Instead of relying on a single fixed storage capacitor, *CapDYN* reconfigures a capacitor bank during charge and discharge. This allows the system to reach usable voltage earlier, reduce the energy left below the load operating voltage after task execution, and expose intermediate energy states to the runtime. The storage subsystem therefore becomes more than an energy buffer: it becomes a hardware–software interface through which the runtime can decide when a task is energy-admissible.

The hardware evaluation validated autonomous charge and discharge reconfiguration on the implemented capacitor-bank prototypes. *CapDYN* reduced leftover energy from 23.67% of the input energy in the fixed-capacitor baseline to 0.16%–5.57% across the measured configurations, while keeping the measured control-infrastructure consumption below 10 μW. The measurements also showed that the choice of capacitor-bank topology matters. The D1 bank scales with lower infrastructural cost but introduces small equalization transients in some transitions, while the D2 bank reduces those switching losses at the cost of more switches and higher overhead as the number of stages increases. This confirms that dynamic capacitance should be dimensioned as a system-level design choice, rather than evaluated only as an isolated storage component.

In the LTspice-based application-level evaluation, the energy-aware runtime translated the hardware energy states into reliable task execution. By mapping profiled tasks to storage stages and releasing each task only when enough energy is available, *CapDYN* avoided failed executions and made useful completions available earlier. In the evaluated workloads, *CapDYN* improved effective efficiency by up to 34.10% with respect to the unaware fixed-capacitor baseline and achieved aggregate first-completion advances of up to 45.00% with respect to the evaluated single-capacitor baselines. The ablation benchmark further clarified the role of hardware–software co-design: the reconfigurable bank alone exposes intermediate energy states, but the runtime is needed to convert those states into earlier task releases. Thus, the main benefit of *CapDYN* is not only higher storage efficiency, but the ability to make harvested en-

ergy available to the application at the right time and in the right amount.

Taken together, the two parts of the thesis show that battery-free design for sustainable IoT systems should be approached as a source-to-load energy-management problem. On the source side, the PMFC interface determines how energy enters the node and how the living electrochemical source is perturbed by the harvesting strategy. On the storage side, *CapDYN* determines how the harvested energy is accumulated, exposed to the runtime, and converted into completed tasks. Both contributions therefore address the same fundamental problem from complementary directions: the PMFC interface increases the opportunity to extract energy from a weak sustainable source, while dynamic capacitance increases the usefulness of the energy once it has been harvested. The source-to-load perspective links these separately evaluated contributions at the system level: future end-to-end PMFC-to-load validation will extend the current source- and storage-side evaluations to the complete integrated chain.

Several limitations remain. The PMFC experiments were performed on a limited number of cells and over a limited observation window compared with the timescales of plant growth, microbial adaptation, electrode aging, and seasonal environmental variation. Larger PMFC populations, longer deployments, and controlled hydration, plant-physiology, and microbiological measurements are needed to separate the effects of intermittent timing from the natural variability of the source. Similarly, the EIS analysis demonstrated measurable loading-history effects, but did not identify their biological origin. On the storage side, the *CapDYN* hardware was evaluated through discrete-component prototypes and application-level LTspice simulations. Future integrated implementations are needed to reduce area, leakage, parasitic losses, and infrastructural consumption, and to validate the architecture in fully autonomous deployments coupled directly to real PMFC energy traces.

Future work should therefore proceed toward a complete PMFC-powered battery-free node in which source control, diagnostic observation, adaptive storage, and task scheduling operate together. On the PMFC side, the intermittent interface can be extended toward closed-loop maximum-power or maximum-energy tracking strategies that adapt the operating voltage, extraction period, and duty cycle using both power measurements and electrochemical indicators.

On the storage side, workload-aware dimensioning algorithms can place capacitance stages near the dominant task-energy clusters, reducing overhead while preserving useful energy granularity. An unresolved system-level question is how the PMFC operating voltage, extraction frequency, and duty cycle set by the intermittent-harvesting interface affect average harvested power and, through *CapDYN*'s charging dynamics, task execution frequency relative to a continuously loaded PMFC supplying the same storage architecture. Finally, integrating the timing interface, PMFC-specific power management, dynamic capacitance, and energy-aware runtime into a single autonomous platform would make it possible to evaluate the complete source-to-load chain under real environmental conditions.

In conclusion, this thesis contributes to the development of sustainable battery-free IoT systems by showing that scarce PMFC energy must be managed jointly at the source, storage, and runtime levels. Controlled intermittent harvesting shows potential to improve the way energy is extracted from a living electrochemical source under suitable operating conditions, while autonomous dynamic capacitance improves the way that energy is buffered and delivered to heterogeneous workloads. The resulting source-to-load perspective provides a foundation for future battery-free systems that are not only powered by sustainable energy sources, but also designed to use every harvested microjoule more effectively.

Appendix A

Appendix

A.1 Post-Switch Voltage and Ideal Switching-Loss Computation

This appendix summarizes the analytical expressions used to compute the voltage immediately after a capacitor-bank reconfiguration event and the corresponding ideal charge-sharing loss. The objective is to support the load-voltage-range discussion in Sec. 4.2 and to clarify the difference between the D1 and D2 bank topologies. Only ideal capacitor-network effects are considered here; switch resistance, leakage, comparator consumption, and other circuit-level parasitics are accounted for separately in the hardware and LTspice evaluations.

A.1.1 General notation

Let V_{SH} and V_{SL} be the upper and lower switching thresholds of the bank, respectively. The stage index is denoted by j, with $j = 1$ corresponding to the minimum-capacitance stage and $j = N$ to the maximum-capacitance stage. The post-switch voltage is the terminal voltage of the capacitor bank immediately after a reconfiguration event and before the next charge or discharge interval.

For a capacitor C_k, let q_k be its stored charge at the switching instant. The voltage across an isolated capacitor is therefore q_k/C_k.

A.1.2 D1 post-switch voltage

In D1, stage j is composed of a parallel block A, containing capacitors $C_1, \ldots, C_j$, in series with a block B, containing capacitors $C_{j+1}, \ldots, C_N$. The post-switch terminal voltage at stage j is the sum of the voltage across the parallel block and the voltages across the series block:

$$V_{\text{post}}^{D1}(j) = \frac{\sum_{k=1}^{j} q_k}{\sum_{k=1}^{j} C_k} + \sum_{k=j+1}^{N} \frac{q_k}{C_k}. \tag{A.1}$$

During charging, the transition from stage j to stage $j+1$ moves capacitor C_{j+1} into the parallel block. After charge equalization within the new parallel block, the post-switch voltage is

$$V_{\text{post},\uparrow}^{D1}(j+1) = \frac{\sum_{k=1}^{j+1} q_k}{\sum_{k=1}^{j+1} C_k} + \sum_{k=j+2}^{N} \frac{q_k}{C_k}. \tag{A.2}$$

During discharge, the transition from stage j to stage $j-1$ occurs after the terminal voltage has reached V_{SL}. Capacitor C_j is removed from the parallel block and inserted into the series block. Assuming no charge redistribution during this topology change, the post-switch voltage is

$$V_{\text{post},\downarrow}^{D1}(j-1) = \frac{\sum_{k=1}^{j-1} q_k}{\sum_{k=1}^{j-1} C_k} + \sum_{k=j}^{N} \frac{q_k}{C_k}. \tag{A.3}$$

Thus, in D1, the post-switch voltage depends on the charge distribution across the individual capacitors. During discharge, the voltage spike is largest near the high-capacitance stages and becomes progressively smaller as the bank moves toward lower-capacitance stages.

A.1.3 D2 post-switch voltage and load-range condition

In D2, the bank is made of 2^{N-1} identical capacitors. At stage j, the capacitors are arranged into 2^{N-j} identical series-connected blocks, each containing 2^{j-1} capacitors in parallel. Because all blocks are identical, the post-switch voltage can be expressed directly.

During charging, the transition from stage j to stage $j+1$ connects pairs of equal-voltage blocks in parallel. The block voltage is unchanged, while the number of series blocks is halved. Therefore, if the transition occurs at V_{SH},

$$V^{D2}_{\text{post},\uparrow}(j+1) = \frac{V_{SH}}{2}. \tag{A.4}$$

During discharge, the transition from stage j to stage $j-1$ splits each parallel block into two equal-voltage sub-blocks and reconnects them in series. The number of series blocks is doubled. Therefore, if the transition occurs at V_{SL},

$$V^{D2}_{\text{post},\downarrow}(j-1) = 2V_{SL}. \tag{A.5}$$

Consequently, D2 produces a fixed halving of the bank voltage during charge-stage increases and a fixed doubling during discharge-stage decreases. To fully exploit all D2 stages without violating the load-voltage range, the upper threshold must therefore satisfy

$$V_{SH} \geq 2V_{SL}. \tag{A.6}$$

If this condition is not met, the full D2 stage range cannot be used without an additional voltage-regulation or smoothing stage.

A.1.4 Ideal charge-sharing loss

The ideal loss associated with a capacitor reconfiguration event can be described as charge-sharing loss. When two capacitive elements C_a and C_b, initially charged at voltages V_a and V_b, are connected in parallel, the energy dissipated during equalization is

$$E_{\text{loss}} = \frac{1}{2}\frac{C_a C_b}{C_a + C_b}\left(V_a - V_b\right)^2. \tag{A.7}$$

This expression shows that ideal redistribution loss occurs only when capacitors or capacitor groups at different voltages are connected in parallel.

For D1, a charging transition from stage j to stage $j+1$ connects capacitor C_{j+1} to the existing parallel block A. Let $C_A(j) = \sum_{k=1}^{j} C_k$, let $V_A(j)$ be the voltage of the parallel block before switching, and let V_{j+1} be the voltage of the incoming capacitor. The ideal charge-sharing loss is

$$E_{\text{loss},\uparrow}^{D1}(j) = \frac{1}{2}\frac{C_A(j)C_{j+1}}{C_A(j)+C_{j+1}}\left(V_A(j)-V_{j+1}\right)^2. \tag{A.8}$$

Therefore, unless the incoming capacitor and the parallel bank are already at the same voltage, D1 exhibits nonzero intrinsic redistribution loss during charge-stage increases. During discharge-stage decreases, the D1 model removes a capacitor from the parallel block and inserts it into the series block without charge equalization; therefore, no ideal parallel charge-sharing loss is produced by the topology change itself.

For D2, charging transitions are obtained by connecting pairs of identical, equal-voltage blocks in parallel. Thus $V_a = V_b$ in Eq. (A.7), and

$$E_{\text{loss},\uparrow}^{D2} = 0. \tag{A.9}$$

Similarly, during discharge, each parallel block is split into two equal-voltage sub-blocks and reconnected in series. No ideal charge redistribution is required, so

$$E_{\text{loss},\downarrow}^{D2} = 0. \tag{A.10}$$

Thus, in the ideal capacitor-only model, D2 has no intrinsic charge-sharing loss at its reconfiguration events, whereas D1 can exhibit redistribution loss during charge-stage increases. At system level, however, this ideal advantage must be balanced against the larger number of capacitors and switches required by D2, which can increase conduction losses, leakage, area, and infrastructural power consumption.

Bibliography

[1] M. Doglioni, E. Yildiz, M. Nardello, K. Akhunov, K. S. Yildirim, and D. Brunelli, "Capdyn: Adaptive self-scaling energy storage for powering batteryless iot," *ACM Trans. Embed. Comput. Syst.*, vol. 24, no. 5, Sep. 2025. [Online]. Available: https://doi.org/10.1145/3737288

[2] M. Nardello, M. Doglioni, S. Dagnino, P. Pastorelli, and D. Brunelli, "Every microjoule counts: Zero-failure task execution in batteryless sensors," in *2025 IEEE SENSORS*, 2025, pp. 1–4.

[3] H. Desai, M. Nardello, D. Brunelli, and B. Lucia, "Camaroptera: A long-range image sensor with local inference for remote sensing applications," *ACM Trans. Embed. Comput. Syst.*, vol. 21, no. 3, May 2022. [Online]. Available: https://doi.org/10.1145/3510850

[4] M. Nardello, L. Caronti, and D. Brunelli, "Intermittent intelligent camera with leo sensor-to-satellite connectivity," in *Proceedings of the 11th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems*, ser. ENSsys '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 79–85. [Online]. Available: https://doi.org/10.1145/3628353.3628550

[5] S. Ahmed, B. Islam, K. S. Yildirim, M. Zimmerling, P. Pawełczak, M. H. Alizai, B. Lucia, L. Mottola, J. Sorber, and J. Hester, "The internet of batteryless things," *Communications of the ACM*, vol. 67, no. 3, pp. 64–73, 2024.

[6] J. Hester and J. Sorber, "The future of sensing is batteryless, intermittent, and awesome," in *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*, 2017, pp. 1–6.

[7] D. Balsamo, A. S. Weddell, G. V. Merrett, B. M. Al-Hashimi, D. Brunelli, and L. Benini, "Hibernus: Sustaining computation during intermittent supply for energy-harvesting systems," *IEEE Embedded Systems Letters*, vol. 7, no. 1, pp. 15–18, 2015.

[8] I. Rusyn, "Role of microbial community and plant species in performance of plant microbial fuel cells," *Renewable and Sustainable Energy Reviews*, vol. 152, p. 111697, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032121009710

[9] M. Doglioni, R. L. Rosa, M. Nardello, and D. Brunelli, "Energy harvesting strategies for plant microbial fuel cells in sustainable iot applications," in *2024 IEEE SENSORS*, 2024, pp. 1–4.

[10] X. Zhang, A. K. Luther, K. Rabaey, and A. Prévoteau, "Periodic polarization duty cycle tunes performance and adhesion of anodic electroactive biofilms," *Bioelectrochemistry*, vol. 155, p. 108581, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1567539423002189

[11] A. Dewan, C. Donovan, D. Heo, and H. Beyenal, "Evaluating the performance of microbial fuel cells powering electronic devices," *Journal of Power Sources*, vol. 195, no. 1, pp. 90–96, 2010. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775309011719

[12] E. J. Gardel, M. E. Nielsen, P. T. J. Grisdela, and P. R. Girguis, "Duty cycling influences current generation in multi-anode environmental microbial fuel cells," *Environmental Science & Technology*, vol. 46, no. 9, pp. 5222–5229, 2012, pMID: 22497491. [Online]. Available: https://doi.org/10.1021/es204622m

[13] F. Grondin, M. Perrier, and B. Tartakovsky, "Microbial fuel cell operation with intermittent connection of the electrical

load," *Journal of Power Sources*, vol. 208, pp. 18–23, 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775312003114

[14] F. Guo, J. T. Babauta, and H. Beyenal, "Impact of intermittent polarization on electrode-respiring geobacter sulfurreducens biofilms," *Journal of Power Sources*, vol. 406, pp. 96–101, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775318311625

[15] A. ter Heijne, M. Pereira, J. Pereira, and T. Sleutels, "Electron storage in electroactive biofilms," *Trends in Biotechnology*, vol. 39, no. 1, p. 34–42, Jan. 2021. [Online]. Available: http://dx.doi.org/10.1016/j.tibtech.2020.06.006

[16] A. Dewan, H. Beyenal, and Z. Lewandowski, "Intermittent energy harvesting improves the performance of microbial fuel cells," *Environmental Science & Technology*, vol. 43, no. 12, pp. 4600–4605, 2009, pMID: 19603683. [Online]. Available: https://doi.org/10.1021/es8037092

[17] L. Hsu, A. Mohamed, P. T. Ha, J. Bloom, T. Ewing, M. Arias-Thode, B. Chadwick, and H. Beyenal, "The influence of energy harvesting strategies on performance and microbial community for sediment microbial fuel cells," *Journal of The Electrochemical Society*, vol. 164, no. 3, p. H3109–H3114, 2017. [Online]. Available: http://dx.doi.org/10.1149/2.0171703jes

[18] H. Wang, J.-D. Park, and Z. Ren, "Active energy harvesting from microbial fuel cells at the maximum power point without using resistors," *Environmental Science & Technology*, vol. 46, no. 9, pp. 5247–5252, 2012, pMID: 22486712. [Online]. Available: https://doi.org/10.1021/es300313d

[19] J.-D. Park and Z. Ren, "Hysteresis controller based maximum power point tracking energy harvesting system for microbial fuel cells," *Journal of Power Sources*, vol. 205, pp. 151–156, 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775312000948

[20] M. Doglioni, M. Nardello, and D. Brunelli, "Plant microbial fuel cells: Energy sources and biosensors for battery-free

smart agriculture," *IEEE Transactions on AgriFood Electronics*, vol. 2, no. 2, pp. 460–470, 2024.

[21] A. Colin, E. Ruppel, and B. Lucia, "A reconfigurable energy storage architecture for energy-harvesting devices," in *Proc. of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems*, 2018, pp. 767–781.

[22] F. Yang, A. S. Thangarajan, S. Michiels, W. Joosen, and D. Hughes, "Morphy: Software defined charge storage for the iot," in *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 248–260. [Online]. Available: https://doi.org/10.1145/3485730.3485947

[23] J. Hester, L. Sitanayah, and J. Sorber, "Tragedy of the coulombs: Federating energy storage for tiny, intermittently-powered sensors," in *Proc. of the 13th ACM Conference on Embedded Networked Sensor Systems*, 2015, pp. 5–16.

[24] E. Yildiz and K. S. Yildirim, "Defragmenting energy storage in batteryless sensing devices," in *Proc. of the 8th International Workshop on Energy Harvesting and Energy-Neutral Sensing Systems*, 2020, pp. 36–42.

[25] M. Doglioni, O. Casas, M. Nardello, and D. Brunelli, "Self-powered plant microbial fuel cell monitoring node with autonomous dynamic capacitance," in *2025 IEEE International Workshop on Metrology for Industry 4.0 & IoT (MetroInd4.0 & IoT)*, 2025, pp. 179–184.

[26] A. Torrisi, M. Doglioni, L. Gemma, A. Albanese, L. Santoro, M. Nardello, and D. Brunelli, "Batteryless soil eis sensor powered by microbial fuel cell," in *Proceedings of SIE 2022*, G. Cocorullo, F. Crupi, and E. Limiti, Eds. Cham: Springer Nature Switzerland, 2023, pp. 277–282.

[27] D. Ayala-Ruiz, A. Castillo Atoche, E. Ruiz-Ibarra, E. Osorio, and J. Castillo, "A self-powered pmfc-based wireless sensor node for smart city applications," *Wireless Communications and Mobile Computing*, vol. 2019, pp. 1–10, 06 2019.

[28] D. Sartori and D. Brunelli, "A smart sensor for precision agriculture powered by microbial fuel cells," in *2016 IEEE Sensors Applications Symposium (SAS)*, 2016, pp. 1–6.

[29] A. Albanese, M. Nardello, G. Fiacco, and D. Brunelli, "Tiny machine learning for high accuracy product quality inspection," *IEEE Sensors Journal*, vol. 23, no. 2, p. 1575 – 1583, 2023.

[30] M. Nardello, H. Desai, D. Brunelli, and B. Lucia, "Camaroptera: a batteryless long-range remote visual sensing system," in *Proceedings of the 7th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems*, ser. ENSsys '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 8–14. [Online]. Available: https://doi.org/10.1145/3362053.3363491

[31] L. Caronti, K. Akhunov, M. Nardello, K. S. Yldlrlm, and D. Brunelli, "Fine-grained hardware acceleration for efficient batteryless intermittent inference on the edge," *ACM Transactions on Embedded Computing Systems*, vol. 22, no. 5, 2023.

[32] D. Brunelli, D. Dondi, A. Bertacchini, L. Larcher, P. Pavan, and L. Benini, "Photovoltaic scavenging systems: Modeling and optimization," *Microelectronics Journal*, vol. 40, no. 9, pp. 1337–1344, 2009.

[33] I. C. Gil-García, A. Fernández-Guillamón, and A. H. Montes-Torres, "Innovation in clean energy from man-made wind and small-wind generation," *Scientific Reports*, vol. 14, no. 1, Oct. 2024. [Online]. Available: http://dx.doi.org/10.1038/s41598-024-74141-w

[34] G. Pasquini, M. Doglioni, A. Saviane, and D. Brunelli, "Self-powered co-harvesting for batteryless sensing and monitoring in two-wheeled iot applications," in *2025 IEEE International Workshop on Sport, Technology and Research (STAR)*, 2025, pp. 239–244.

[35] J.-L. Unuofin, S. Iwarere, and M. Daramola, "Embracing the future of circular bio-enabled economy: unveiling the

prospects of microbial fuel cells in achieving true sustainable energy," *Environmental Science and Pollution Research*, vol. 30, pp. 1–27, 07 2023.

[36] V. Ancona, C. Cavone, P. Grenni, G. Gagliardi, C. Cosentini, D. Borello, and A. Barra Caracciolo, "Plant microbial fuel cells for recovering contaminated environments," *International Journal of Hydrogen Energy*, vol. 72, pp. 1116–1126, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360319924021670

[37] P. Jyoti Sarma and K. Mohanty, "A novel three-chamber modular pmfc with bentonite/flyash based clay membrane and oxygen reducing biocathode for long term sustainable bioelectricity generation," *Bioelectrochemistry*, vol. 144, p. 107996, 2022.

[38] K. Wetser, E. Sudirjo, C. J. Buisman, and D. P. Strik, "Electricity generation by a plant microbial fuel cell with an integrated oxygen reducing biocathode," *Applied Energy*, vol. 137, pp. 151–157, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261914010460

[39] R. Agrahari, B. Bayar, H. N. Abubackar, B. S. Giri, E. R. Rene, and R. Rani, "Advances in the development of electrode materials for improving the reactor kinetics in microbial fuel cells," *Chemosphere*, vol. 290, p. 133184, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0045653521036560

[40] R. Piyare, A. L. Murphy, P. Tosato, and D. Brunelli, "Plug into a plant: Using a plant microbial fuel cell and a wake-up radio for an energy neutral sensing system," in *2017 IEEE 42nd Conference on Local Computer Networks Workshops (LCN Workshops)*, 2017, pp. 18–25.

[41] A. Rodriguez Arreola, D. Balsamo, A. K. Das, A. S. Weddell, D. Brunelli, B. M. Al-Hashimi, and G. V. Merrett, "Approaches to transient computing for energy harvesting systems: A quantitative evaluation," in *Proceedings of the 3rd International Workshop on Energy Harvesting & Energy Neutral Sensing Systems*, ser. ENSsys '15. New York,

NY, USA: Association for Computing Machinery, 2015, p. 3–8. [Online]. Available: https://doi.org/10.1145/2820645.2820652

[42] F. T. Kabutey, Q. Zhao, L. Wei, J. Ding, P. Antwi, F. K. Quashie, and W. Wang, "An overview of plant microbial fuel cells (pmfcs): Configurations and applications," *Renewable and Sustainable Energy Reviews*, vol. 110, pp. 402–414, 2019.

[43] P. de Jager, D. Groen, and D. P. Strik, "Long-term performance of pilot-scale tubular plant-microbial fuel cells in a brownfield-constructed wetland," *Renewable Energy*, vol. 219, p. 119532, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0960148123014477

[44] S. Maye, L. Delabays, J. Sansonnens, M. Blatter, G. Huguenin, and F. Fischer, "Computerized voltage reversal prevention in second and third year 1000-liter microbial fuel cell," *Renewable and Sustainable Energy Reviews*, vol. 208, p. 115017, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032124007433

[45] B. Kim, S. V. Mohan, D. Fapyane, and I. S. Chang, "Controlling voltage reversal in microbial fuel cells," *Trends in Biotechnology*, vol. 38, no. 6, pp. 667–678, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167779919303026

[46] H. Xu, X.-L. Yang, Z.-H. Zhang, Y.-G. Xia, and H.-L. Song, "External circuit loading mode regulates anode biofilm electrochemistry and pollutants removal in microbial fuel cells," *Bioresource Technology*, vol. 410, p. 131300, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0960852424010046

[47] P. Bhavya, V. Sonali, M. Prerna, and A. Sandeep, "Electrochemical impedance spectroscopy (eis) performance analysis and challenges in fuel cell applications," *J. Electrochem. Sci. Technol*, vol. 13, no. 2, pp. 167–176, 2022. [Online]. Available: http://www.jecst.org/journal/view.php?number=436

[48] H. Wang, X. Long, Y. Sun, D. Wang, Z. Wang, H. Meng, C. Jiang, W. Dong, and N. Lu, "Electrochemical impedance spectroscopy applied to microbial fuel cells: A review," *Frontiers in Microbiology*, vol. Volume 13 - 2022, 2022. [Online]. Available: https://www.frontiersin.org/journals/microbiology/articles/10.3389/fmicb.2022.973501

[49] M. Rossi, A. Iannaci, P. Tosato, and D. Brunelli, "Let the microbes power your sensing display," in *2017 IEEE SENSORS*, 2017, pp. 1–3.

[50] A. Costilla Reyes, C. Erbay, S. Carreon-Bautista, A. Han, and E. Sánchez-Sinencio, "A time-interleave-based power management system with maximum power extraction and health protection algorithm for multiple microbial fuel cells for internet of things smart nodes," *Applied Sciences*, vol. 8, no. 12, 2018. [Online]. Available: https://www.mdpi.com/2076-3417/8/12/2404

[51] A. Mohamed, H. M. Zmuda, P. T. Ha, E. R. Coats, and H. Beyenal, "Large-scale switchable potentiostatically controlled/microbial fuel cell bioelectrochemical wastewater treatment system," *Bioelectrochemistry*, vol. 138, p. 107724, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1567539420306472

[52] R. Rossi, B. P. Cario, C. Santoro, W. Yang, P. E. Saikaly, and B. E. Logan, "Evaluation of electrode and solution area-based resistances enables quantitative comparisons of factors impacting microbial fuel cell performance," *Environmental Science &; Technology*, vol. 53, no. 7, p. 3977–3986, Feb. 2019. [Online]. Available: http://dx.doi.org/10.1021/acs.est.8b06004

[53] H.-S. Lee, C. I. Torres, and B. E. Rittmann, "Effects of substrate diffusion and anode potential on kinetic parameters for anode-respiring bacteria," *Environmental Science &; Technology*, vol. 43, no. 19, p. 7571–7577, 2009. [Online]. Available: http://dx.doi.org/10.1021/es9015519

[54] S. Jung, M. M. Mench, and J. M. Regan, "Impedance characteristics and polarization behavior of a microbial

fuel cell in response to short-term changes in medium ph," *Environmental Science &; Technology*, vol. 45, no. 20, p. 9069–9074, 2011. [Online]. Available: http://dx.doi.org/10.1021/es201737g

[55] X. Zhang, A. Prévoteau, R. O. Louro, C. M. Paquete, and K. Rabaey, "Periodic polarization of electroactive biofilms increases current density and charge carriers concentration while modifying biofilm structure," *Biosensors and Bioelectronics*, vol. 121, p. 183–191, Dec. 2018. [Online]. Available: http://dx.doi.org/10.1016/j.bios.2018.08.045

[56] N. Uría, X. Muñoz Berbel, O. Sánchez, F. X. Muñoz, and J. Mas, "Transient storage of electrical charge in biofilms of ¡i¿shewanella oneidensis¡/i¿ mr-1 growing in a microbial fuel cell," *Environmental Science &; Technology*, vol. 45, no. 23, p. 10250–10256, Oct. 2011. [Online]. Available: http://dx.doi.org/10.1021/es2025214

[57] L. Caizán-Juanarena, C. Borsje, T. Sleutels, D. Yntema, C. Santoro, I. Ieropoulos, F. Soavi, and A. ter Heijne, "Combination of bioelectrochemical systems and electrochemical capacitors: Principles, analysis and opportunities," *Biotechnology Advances*, vol. 39, p. 107456, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0734975019301569

[58] R. Fernandez Feito, T. Younas, and R. M. Dinsdale, "Evaluation of a comprehensive power management system with maximum power point tracking algorithm for multiple microbial fuel cell energy harvesting," *Bioelectrochemistry*, vol. 155, p. 108597, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1567539423002347

[59] D. Sun, S. Cheng, F. Zhang, and B. E. Logan, "Current density reversibly alters metabolic spatial structure of exoelectrogenic anode biofilms," *Journal of Power Sources*, vol. 356, pp. 566–571, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775316316822

[60] M. Fernandez-Gatell, X. Sanchez-Vila, and J. Puigagut, "Exploring the biocapacitance in m3c-based biosensors for

the assessment of microbial activity and organic matter," *Science of The Total Environment*, vol. 904, p. 166510, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0048969723051355

[61] T. H. Chung, M. N. Meshref, and B. R. Dhar, "Microbial electrochemical biosensor for rapid detection of naphthenic acid in aqueous solution," *Journal of Electroanalytical Chemistry*, vol. 873, p. 114405, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1572665720306329

[62] P. Liang, C. Zhang, Y. Jiang, Y. Bian, H. Zhang, X. Sun, X. Yang, X. Zhang, and X. Huang, "Performance enhancement of microbial fuel cell by applying transient-state regulation," *Applied Energy*, vol. 185, pp. 582–588, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S030626191631580X

[63] J. Coronado, M. Perrier, and B. Tartakovsky, "Pulse-width modulated external resistance increases the microbial fuel cell power output," *Bioresource Technology*, vol. 147, pp. 65–70, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0960852413012340

[64] H. Wang, Z. Ren, and J.-D. Park, "Power electronic converters for microbial fuel cell energy extraction: Effects of inductance, duty ratio, and switching frequency," *Journal of Power Sources*, vol. 220, pp. 89–94, 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S037877531201227X

[65] M. Quaglio, D. Ahmed, G. Massaglia, A. Sacco, V. Margaria, and C. F. Pirri, "Enhanced power extraction with sediment microbial fuel cells by anode alternation," *Fuels*, vol. 2, no. 2, pp. 168–178, 2021. [Online]. Available: https://www.mdpi.com/2673-3994/2/2/10

[66] Y. Kim, M. C. Hatzell, A. J. Hutchinson, and B. E. Logan, "Capturing power at higher voltages from arrays of microbial fuel cells without voltage reversal," *Energy Environ. Sci.*, vol. 4, pp. 4662–4667, 2011. [Online]. Available: http://dx.doi.org/10.1039/C1EE02451E

[67] E. Osorio-de-la Rosa, J. Vázquez-Castillo, A. Castillo-Atoche, J. Heredia-Lozano, A. Castillo-Atoche, G. Becerra-Nuñez, and R. Barbosa, "Arrays of plant microbial fuel cells for implementing self-sustainable wireless sensor networks," *IEEE Sensors Journal*, vol. 21, no. 2, pp. 1965–1974, 2021.

[68] S. Feng, B. Babaiahgari, T. Kim, and J.-D. Park, "Modular multilevel converter-based energy harvester for microbial fuel cells stacked in series," *Journal of Power Sources*, vol. 656, p. 237954, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775325017902

[69] J. An, J. Sim, and H.-S. Lee, "Control of voltage reversal in serially stacked microbial fuel cells through manipulating current: Significance of critical current density," *Journal of Power Sources*, vol. 283, pp. 19–23, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775315003109

[70] M. Alaraj and J.-D. Park, "Net power positive maximum power point tracking energy harvesting system for microbial fuel cell," *Journal of Power Sources*, vol. 418, pp. 225–232, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775319301594

[71] M. Alaraj, M. Radenkovic, and J.-D. Park, "Intelligent energy harvesting scheme for microbial fuel cells: Maximum power point tracking and voltage overshoot avoidance," *Journal of Power Sources*, vol. 342, pp. 726–732, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775316318079

[72] S. M. de Smit, C. J. N. Buisman, J. H. Bitter, and D. P. B. T. B. Strik, "Cyclic voltammetry is invasive on microbial electrosynthesis," *ChemElectroChem*, vol. 8, no. 17, pp. 3384–3396, 2021. [Online]. Available: https://chemistry-europe.onlinelibrary.wiley.com/doi/abs/10.1002/celc.202100914

[73] F. Ma, Y. Yin, and M. Li, "Start-up process modelling of sediment microbial fuel cells based on data driven," *Mathematical Problems in Engineering*, vol. 2019, pp. 1–10, 01 2019.

[74] M. Sindhuja, N. S. Kumar, V. Sudha, and S. Harinipriya, "Equivalent circuit modeling of microbial fuel cells using impedance spectroscopy," *Journal of Energy Storage*, vol. 7, pp. 136–146, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352152X16300858

[75] A. C. Sophia and S. Sreeja, "Green energy generation from plant microbial fuel cells (pmfc) using compost and a novel clay separator," *Sustainable Energy Technologies and Assessments*, vol. 21, pp. 59–66, 2017.

[76] P. Si, Y. Feng, X. Wang, X. Yue, D. Li, S. Zhang, B. Liang, Y. Gao, and A. Zhou, "Enhanced degradation of metronidazole by the coupling of photocatalytic and microbial fuel cell: Mechanism and electrochemistry characteristic," *Journal of Environmental Chemical Engineering*, vol. 11, no. 3, p. 109707, 2023.

[77] D. Brunelli, P. Tosato, and M. Rossi, "Microbial fuel cell as a biosensor and a power source for flora health monitoring," in *2016 IEEE SENSORS*, 2016, pp. 1–3.

[78] B. Koo, H. Chai, S. Son, N. Islam, C. Hong, and S. P. Jung, "Potential-dependent impedance characteristics of bioanodes and performance enhancement using phosphoric acid-treated anodes in continuous-flow microbial fuel cells," *Chemical Engineering Journal*, vol. 519, p. 165285, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1385894725061212

[79] K. ARIYOSHI, Z. SIROMA, A. MINESHIGE, M. TAKENO, T. FUKUTSUKA, T. ABE, and S. UCHIDA, "Electrochemical impedance spectroscopy part 1: Fundamentals," *Electrochemistry*, vol. 90, no. 10, p. 102007–102007, Oct. 2022. [Online]. Available: http://dx.doi.org/10.5796/electrochemistry.22-66071

[80] A. C. Lazanas and M. I. Prodromidis, "Electrochemical impedance spectroscopy—a tutorial," *ACS Measurement Science Au*, vol. 3, no. 3, p. 162–193, Mar. 2023. [Online]. Available: http://dx.doi.org/10.1021/acsmeasuresciau.2c00070

[81] Z. He and F. Mansfeld, "Exploring the use of electrochemical impedance spectroscopy (eis) in microbial fuel cell studies," *Energy Environ. Sci.*, vol. 2, no. 2, p. 215–219, 2009. [Online]. Available: http://dx.doi.org/10.1039/b814914c

[82] M. Sindhuja, N. S. Kumar, V. Sudha, and S. Harinipriya, "Equivalent circuit modeling of microbial fuel cells using impedance spectroscopy," *Journal of Energy Storage*, vol. 7, p. 136–146, Aug. 2016. [Online]. Available: http://dx.doi.org/10.1016/j.est.2016.06.005

[83] N. S. Ramaraja P Ramasamy, "Electrochemical impedance spectroscopy for microbial fuel cell characterization," *Journal of Microbial & Biochemical Technology*, 2013. [Online]. Available: http://dx.doi.org/10.4172/1948-5948.s6-004

[84] R. P. Ramasamy, Z. Ren, M. M. Mench, and J. M. Regan, "Impact of initial biofilm growth on the anode impedance of microbial fuel cells," *Biotechnology and Bioengineering*, vol. 101, no. 1, pp. 101–108, 2008. [Online]. Available: https://analyticalsciencejournals.onlinelibrary.wiley.com/doi/abs/10.1002/bit.21878

[85] Z. Lu, P. Girguis, P. Liang, H. Shi, G. Huang, L. Cai, and L. Zhang, "Biological capacitance studies of anodes in microbial fuel cells using electrochemical impedance spectroscopy," *Bioprocess and Biosystems Engineering*, vol. 38, no. 7, p. 1325–1333, Feb. 2015. [Online]. Available: http://dx.doi.org/10.1007/s00449-015-1373-z

[86] A. ter Heijne, D. Liu, M. Sulonen, T. Sleutels, and F. Fabregat-Santiago, "Quantification of bio-anode capacitance in bioelectrochemical systems using electrochemical impedance spectroscopy," *Journal of Power Sources*, vol. 400, pp. 533–538, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775318308620

[87] A. Sabovic, A. K. Sultania, C. Delgado, L. D. Roeck, and J. Famaey, "An energy-aware task scheduler for energy-harvesting batteryless iot devices," *IEEE Internet of Things Journal*, vol. 9, no. 22, pp. 23 097–23 114, 2022.

[88] V. Gogolou, S. Siskos, and T. Laopoulos, "A reconfigurable storage unit for autonomously powered iot nodes," in *2019 Panhellenic Conference on Electronics & Telecommunications (PACET)*, 2019, pp. 1–4.

[89] F. El Mahboubi, M. Bafleur, V. Boitier, and J.-M. Dilhac, "Energy-harvesting powered variable storage topology for battery-free wireless sensors," *Technologies*, vol. 6, no. 4, 2018.

[90] S. Mileiko, O. Cetinkaya, D. Mackie, A. Yakovlev, and D. Balsamo, "A non-volatile state retention unit for multi-storage energy management in transient systems," in *2023 9th International Workshop on Advances in Sensors and Interfaces (IWASI)*, 2023, pp. 46–51.

[91] J. Hester and J. Sorber, "Flicker: Rapid prototyping for the batteryless internet-of-things," in *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*, 2017, pp. 1–13.

[92] A. Alsubhi, S. Babatunde, N. Tobias, and J. Sorber, "Stash: Flexible energy storage for intermittent sensors," *ACM Trans. Embed. Comput. Syst.*, jan 2024, just Accepted. [Online]. Available: https://doi.org/10.1145/3641511

[93] M. Doglioni, "CapDYN: Thesis artifacts," 2026. [Online]. Available: https://doi.org/10.5281/zenodo.20747777

[94] A. Siskos, F. E. Mahboubi, V. Boitier, T. Laopoulos, and M. Bafleur, "A power management system using reconfigurable storage scheme for batteryless wireless sensor nodes," in *2018 7th International Conference on Modern Circuits and Systems Technologies (MOCAST)*, 2018, pp. 1–4.

[95] A. Bakar, R. Goel, J. de Winkel, J. Huang, S. Ahmed, B. Islam, P. Pawełczak, K. S. Yıldırım, and J. Hester, "Protean: An energy-efficient and heterogeneous platform for adaptive and hardware-accelerated battery-free computing," in *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, 2022, pp. 207–221.

[96] J. de Winkel, V. Kortbeek, J. Hester, and P. Pawelczak, "Battery-free game boy," *Proceedings of the ACM on Inter-*

active, Mobile, Wearable and Ubiquitous Technologies, vol. 4, no. 111, pp. 1–26, 2020.

[97] K. Maeng and B. Lucia, "Adaptive low-overhead scheduling for periodic and reactive intermittent execution," in *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2020. New York, NY, USA: Association for Computing Machinery, 2020, p. 1005–1021. [Online]. Available: https://doi.org/10.1145/3385412.3385998

[98] E. Yildiz, S. Ahmed, B. Islam, J. Hester, and K. S. Yildirim, "Efficient and safe i/o operations for intermittent systems," in *Proceedings of the Eighteenth European Conference on Computer Systems*, 2023, pp. 63–78.

[99] E. Yıldız, L. Chen, and K. S. Yıldırım, "Immortal threads: Multithreaded event-driven intermittent computing on {Ultra-Low-Power} microcontrollers," in *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, 2022, pp. 339–355.

[100] J. de Winkel, C. Delle Donne, K. S. Yildirim, P. Pawełczak, and J. Hester, "Reliable timekeeping for intermittent computing," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 53–67. [Online]. Available: https://doi.org/10.1145/3373376.3378464

[101] P. Singla and S. R. Sarangi, "A survey and experimental analysis of checkpointing techniques for energy harvesting devices," *Journal of Systems Architecture*, vol. 126, p. 102464, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1383762122000534

[102] N. A. Bhatti and L. Mottola, "Harvos: efficient code instrumentation for transiently-powered embedded sensing," in *Proceedings of the 16th ACM/IEEE International Conference on Information Processing in Sensor Networks*, ser. IPSN '17. New York, NY, USA: Association for

Computing Machinery, 2017, p. 209–219. [Online]. Available: https://doi.org/10.1145/3055031.3055082

[103] D. Balsamo, A. S. Weddell, A. Das, A. R. Arreola, D. Brunelli, B. M. Al-Hashimi, G. V. Merrett, and L. Benini, "Hibernus++: a self-calibrating and adaptive system for transiently-powered embedded devices," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 35, no. 12, pp. 1968–1980, 2016.

[104] K. Maeng and B. Lucia, "Adaptive dynamic checkpointing for safe efficient intermittent computing," in *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, 2018, pp. 129–144.

[105] Y. Kim, Y. Lim, and C. Lim, "Liveness-aware checkpointing of arrays for efficient intermittent computing," in *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2023, pp. 1–6.

[106] S. Ahmed, N. A. Bhatti, M. H. Alizai, J. H. Siddiqui, and L. Mottola, "Efficient intermittent computing with differential checkpointing," in *Proceedings of the 20th ACM SIGPLAN/SIGBED International Conference on Languages, Compilers, and Tools for Embedded Systems*, 2019, pp. 70–81.

[107] A. Colin and B. Lucia, "Chain: Tasks and channels for reliable intermittent programs," in *Proc. OOPSLA*. Amsterdam, Netherlands: ACM, 2016, pp. 514–530.

[108] K. S. Yıldırım, A. Y. Majid, D. Patoukas, K. Schaper, P. Pawelczak, and J. Hester, "Ink: Reactive kernel for tiny batteryless sensors," in *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*, 2018, pp. 41–53.

[109] K. Maeng, A. Colin, and B. Lucia, "Alpaca: Intermittent execution without checkpoints," *Proceedings of the ACM on Programming Languages*, vol. 1, no. OOPSLA, pp. 1–30, 2017.

[110] S. Liu, W. Zhang, M. Lv, Q. Chen, and N. Guan, "Latics: A low-overhead adaptive task-based intermittent computing system," *IEEE Transactions on Computer-Aided Design of*

Integrated Circuits and Systems, vol. 39, no. 11, pp. 3711–3723, 2020.

[111] A. Y. Majid, C. D. Donne, K. Maeng, A. Colin, K. S. Yildirim, B. Lucia, and P. Pawełczak, "Dynamic task-based intermittent execution for energy-harvesting devices," *ACM Transactions on Sensor Networks (TOSN)*, vol. 16, no. 1, pp. 1–24, 2020.

[112] M. Karimi and H. Kim, "Energy scheduling for task execution on intermittently-powered devices," *SIGBED Rev.*, vol. 17, no. 1, p. 36–41, Jul. 2020. [Online]. Available: https://doi.org/10.1145/3412821.3412827

[113] F. Erata, E. Yildiz, A. Goknil, K. S. Yildirim, J. Szefer, R. Piskac, and G. Sezgin, "Etap: Energy-aware timing analysis of intermittent programs," *ACM Trans. Embed. Comput. Syst.*, vol. 22, no. 2, Jan. 2023. [Online]. Available: https://doi.org/10.1145/3563216

[114] C. Delgado and J. Famaey, "Optimal energy-aware task scheduling for batteryless iot devices," *IEEE Transactions on Emerging Topics in Computing*, vol. 10, no. 3, pp. 1374–1387, 2022.

[115] E. Ruppel, M. Surbatovich, H. Desai, K. Maeng, and B. Lucia, "An architectural charge management interface for energy-harvesting systems," in *2022 55th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2022, pp. 318–335.

[116] J. Choi, H. Joe, and C. Jung, "Capos: Capacitor error resilience for energy harvesting systems," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 11, pp. 4539–4550, 2022.

[117] Texas Instruments, "Msp430fr58xx, msp430fr59xx, msp430fr68xx, and msp430fr69xx family user's guide," http://www.ti.com/lit/ug/slau367o/slau367o.pdf, 2019, last accessed: September 2019.

[118] J. Hester, T. Scott, and J. Sorber, "Ekho: Realistic and repeatable experimentation for tiny energy-harvesting sensors,"

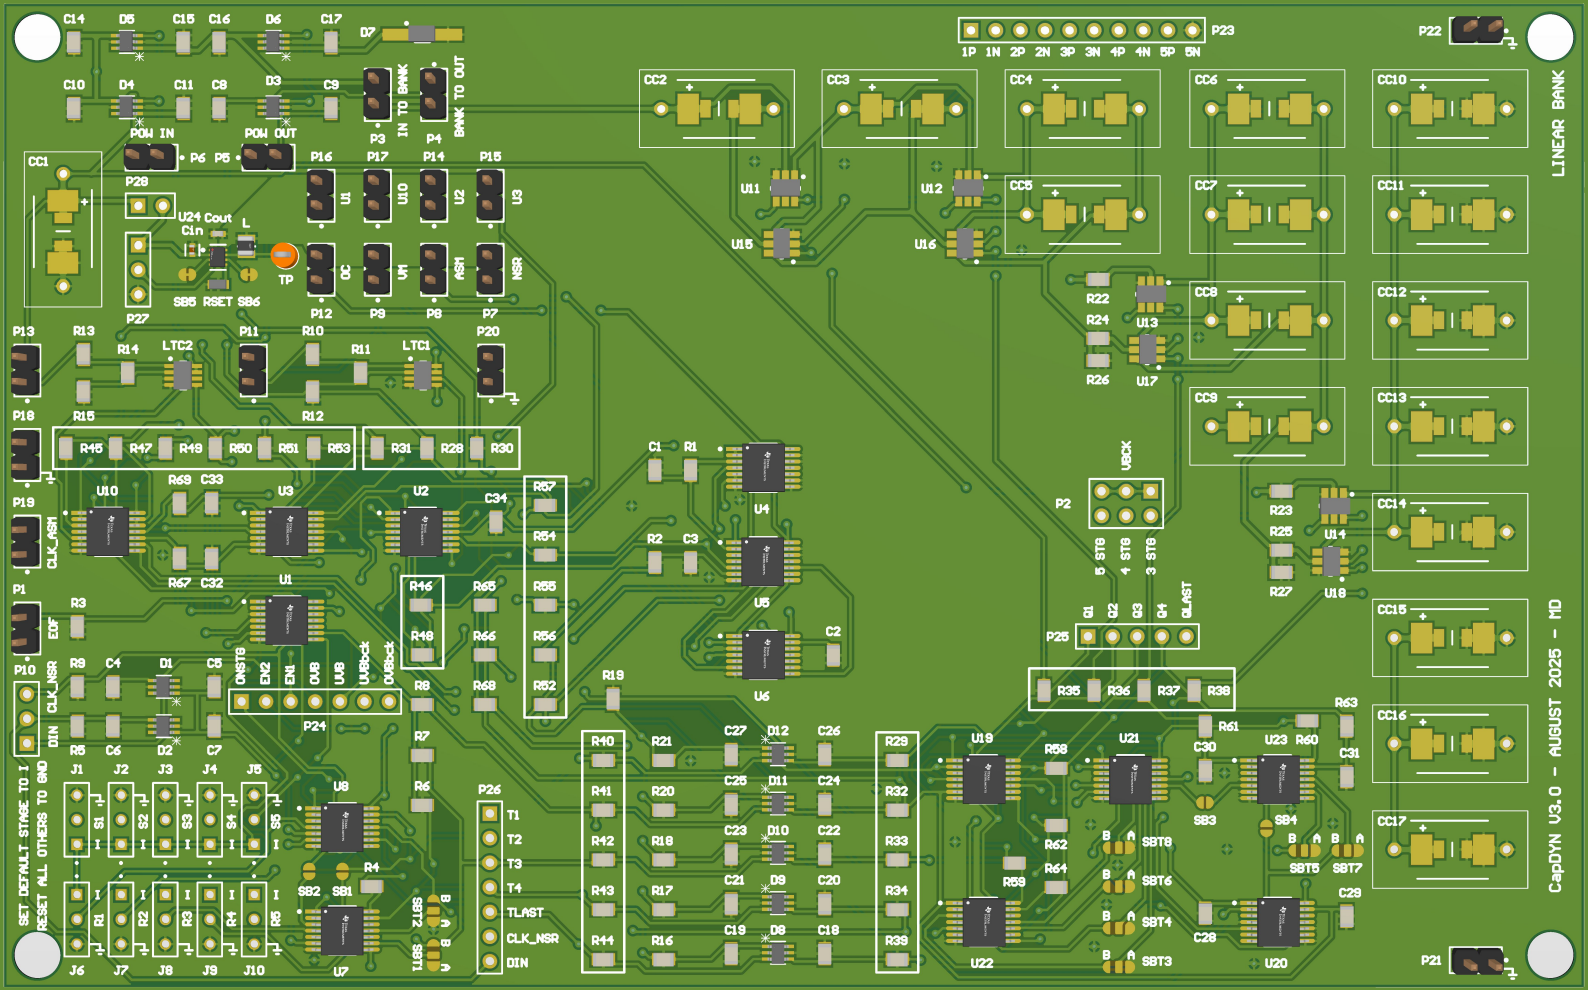
in *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*, 2014, pp. 330–331.

[119] K. Y. Cheng, R. Cord-Ruwisch, and G. Ho, "A new approach for in situ cyclic voltammetry of a microbial fuel cell biofilm without using a potentiostat," *Bioelectrochemistry*, vol. 74, no. 2, pp. 227–231, 2009. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1567539408001722

[120] P. Ibba, M. Crepaldi, G. Cantarella, G. Zini, A. Barcellona, M. Petrelli, B. D. Abera, B. Shkodra, L. Petti, and P. Lugli, "Fruitmeter: An ad5933-based portable impedance analyzer for fruit quality characterization," in *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2020, pp. 1–5.

[121] M. Gorlatova, M. Zapas, E. Xu, M. Bahlke, I. J. Kymissis, and G. Zussman, "Crawdad columbia/enhants," 2022. [Online]. Available: https://dx.doi.org/10.15783/C77P47

[122] S. Kurpaska, J. Knaga, H. Latala, J. Sikora, and W. Tomczyk, "Efficiency of solar radiation conversion in photovoltaic panels," *BIO Web of Conferences*, vol. 10, p. 02014, 01 2018.


LINEAR BANK
CapDYN V3.0 - AUGUST 2025 - MD
SET DEFAULT STAGE TO I
RESET ALL OTHERS TO GND
P23
1P 1N 2P 2N 3P 3N 4P 4N 5P 5N
P22
P21
P2
VBCK
5 ST6
4 ST6
3 ST6
P25
Q1 Q2 Q3 Q4 QLAST
P26
T1
T2
T3
T4
TLAST
CLK_NSR
DIN
P24
ONST6
EN2
EN1
OV8
UV8
UVBack
OVBack
POW IN
POW OUT
IN TO BANK
BANK TO OUT
P3
P4
P5
P6
P16
P17
P14
P15
U1
U10
U2
U3
P12
P9
P8
P7
OC
UV
ASM
NSR
P28
P27
U24 Cout
Cin
L
SB5
RSET
SB6
TP
P13
P18
P19
CLK_ASM
P1
EXT
P10
CLK_NSR
DIN
P11
P20
LTC2
LTC1
R13
R14
R10
R11
R15
R12
R45
R47
R49
R50
R51
R53
R31
R28
R30
U10
U3
U2
U1
U4
U5
U6
U7
U8
U11
U12
U15
U16
U13
U17
U14
U18
U19
U21
U22
U23
U20
CC1
CC2
CC3
CC4
CC5
CC6
CC7
CC8
CC9
CC10
CC11
CC12
CC13
CC14
CC15
CC16
CC17
SB1
SB2
SB3
SB4
SB5
SB6
SB7
SB8
J1
J2
J3
J4
J5
J6
J7
J8
J9
J10
S1
S2
S3
S4
S5
R1
R2
R3
R4
R5
TLAST