A Hardware-Based Multi-Stage Dynamic Power Management Architecture for Autonomous Low-Light Operation

Charalampos S. Kouzinopoulos
Maastricht University
Maastricht, The Netherlands

Marcel Louis Meli
ZHAW School of Engineering, Zurich University of Applied Sciences
Zurich, Switzerland

Martin Schellenberg
Philip J. Poole
Microdul AG
Zurich, Switzerland

Mathieu Bellanger
Matthias Kauer
Lightricity Limited
Oxford, United Kingdom

Julien De Vos
e-peas S.A.
Mont-Saint-Guibert, Belgium

Dimosthenis Ioannidis
Dimitrios Tzovaras
CERTH/ITI
Thermi, Greece

Abstract

The advance of autonomous Smart Sensor Networks and embedded systems for the Internet of Things, powered by photovoltaic energy harvesting, is severely limited by energy efficiency, especially in low-light environments. While Dynamic Power Management is essential for energy conservation, conventional software-based techniques that rely on processor-managed low-power states incur a persistent quiescent current drain. This current becomes the dominant energy sink in energy-scarce conditions, limiting autonomy. The work of this paper addresses this limitation by introducing a robust, hardware-orchestrated dynamic power management architecture that improves existing configurations for battery-based sensor nodes. The proposed architecture achieves a minimal quiescent drain of $452nA$, by completely power-gating the microcontroller and all non-essential peripherals, with wake-up orchestrated by an ultra-low-power PMIC, RTC and a novel latch circuit developed specifically for this work. Our evaluation demonstrates that the dynamic power management architecture is significantly more efficient than traditional software-based sleep modes.

CCS Concepts

• **Computer systems organization → Embedded systems**; • **Hardware → Power estimation and optimization**; **Sensor devices and platforms**.

Keywords

Low-Power, Power Management, PMIC, Energy Harvesting, Embedded Systems, Internet of Things

ACM Reference Format:

Charalampos S. Kouzinopoulos, Marcel Louis Meli, Martin Schellenberg, Philip J. Poole, Mathieu Bellanger, Matthias Kauer, Julien De Vos, Dimosthenis Ioannidis, and Dimitrios Tzovaras. 2026. A Hardware-Based Multi-Stage Dynamic Power Management Architecture for Autonomous Low-Light Operation. In *Proceedings of IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED 2026).* ACM, New York, NY, USA, 7 pages. https://doi.org/XXXXXXX.XXXXXXX

ISLPED 2026, Chicago, USA

ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

The rapid development of Internet of Things (IoT) applications has led to the widespread deployment of Smart Sensor Nodes (SSNs), embedded, autonomous devices for data-driven automation at the edge. The nodes integrate sensors for data acquisition with wireless communication capabilities, while recent advances have enabled the integration of *edge AI*; the execution of AI models at the edge [2][5][6] for localized decision-making. To achieve autonomy, these systems are often powered by energy harvesting sources such as photovoltaic (PV) elements. In such energy-constrained settings, particularly in indoor or variable outdoor environments, energy budget is largely dictated by low-light performance.

Dynamic power management (DPM) [3][4][11] is a fundamental approach to maximize operational lifetime, typically involving Dynamic Voltage and Frequency Scaling (DVFS), clock gating, establishment of Voltage-Frequency Islands with dynamically adjustable Voltage/Frequency knobs, as well as event-driven sleep and low-power modes. However, these strategies share a fundamental limitation: they are typically software-orchestrated. The main system microcontroller (MCU) must remain in a deep-sleep or standby state to manage wake-up events. This active-sleep state, while low-power, incurs a persistent quiescent current (indicatively $2.57\mu A$ for an STM32L4 MCU in stop 2 mode). In ultra-low-light conditions where harvested power is minimal, this quiescent drain becomes the dominant energy sink, creating an operational floor for autonomous operation.

To improve upon this gap, we build upon and extend recent work in battery-less systems [9][10] to propose a hardware-orchestrated DPM architecture for battery-based sensor nodes that completely power-gates the MCU and most peripherals during sleep cycles, eliminating their quiescent current. This architecture is built around the always-on power domain of a Power Management Integrated Circuit (PMIC), powering only a low-power Real-Time clock (RTC) and a capacitive touch sensor for periodic or on-demand state switching. These hardware components are responsible for orchestrating the wake-up and shutdown of the main compute domain.

Contributions. The main contributions of this work are:

- A robust hardware-orchestrated DPM architecture for battery-based sensor nodes that eliminates MCU and peripheral quiescent current by fully power-gating the compute domain, achieving a measured sleep-state drain of $452nA$.
- A new latch circuit for the capacitive sensor, forming a key part of the DPM architecture. This circuit, which enables the sensor to provide a persistent *enable* signal from a momentary trigger, was designed by the authors and integrated into the commercial Microdul MS8892 capacitive sensor.
- The design of an autonomous, energy-harvesting SSN, integrating the proposed DPM architecture.
- Experimental results demonstrating cold-starting capability and sustained, net-positive energy operation under very low illuminance levels.

2 Related Work

Most prior work on power management optimization for energy-harvesting autonomous SSNs focuses on aggressive duty-cycling, where components enter idle, stop or deep sleep during periods of inactivity. While effective, these approaches still incur significant quiescent currents, limiting energy autonomy under low-light conditions (i.e. [1][12]). To further reduce leakages, several designs actively cut-off power to non-essential components during inactive periods (i.e. [7][8]). However, in these approaches the MCU typically remains powered, reducing system autonomy.

Some works fully power down the main load, including the MCU and most non-essential peripherals, while keeping only a minimal set of ultra-low-power components active [9][10]. These designs achieve significant energy efficiency under ultra-low-light conditions as low as 5-20 lux. However, they are typically battery-less and rely exclusively on small capacitive storage, which severely limits the maximum supported load and edge processing capabilities.

For example, [10] presents a Bluetooth Low Energy (BLE) node operating down to 5 lux by duty-cycling between a 310nA Deep Sleep state and a Normal mode either periodically using an RTC or based on a user command via a capacitive touch sensor. Unlike our proposed architecture though, this system lacks a dedicated PMIC and battery. Instead, it relies solely on a $100\mu F$ capacitor charged directly from PV, restricting its application to minimal, highly intermittent loads. Although a custom ASIC PMIC is included, with a reported operating current of $30nA$ at room temperature, only a single comparator of the Integrated Circuit (IC) is utilized to monitor the voltage of the capacitor; the system can be powered up as long as the supercapacitor's voltage is at least $2.25V$. The quiescent current consumption of the system in Deep Sleep can be approximated to $140nA$ based on the consumption of the comparator, RTC and touch sensor.

Similarly, [9] reports a LoRaWAN node operating down to 20 lux using a dedicated PMIC and a $10mF$ supercapacitor to power a LoRa SX1276 transceiver, a BME680 sensor and a Cortex-M4 MCU, requiring $8mJ$ per 10-minute active cycle. The limits of the PMIC are set in such a way that the system will start operating as soon as the charge of the supercapacitor reaches $3V$ and it will continue charging up to a maximum of $3.6V$. After the $3V$ threshold is reached, the supercapacitor can be discharged down to $2V$. A total load of up to $44.8mJ$ can be supported by this configuration. A larger supercapacitor can be used to support increased loads, at the disadvantage of requiring additional time to charge it at similar illumination levels.

Table 1: State-of-the-art comparison

Work	Always-on components	Always-on current
[10]	MA198 (Comparator), RV-3028-C7, MS8892	$30nA + 45nA + 65nA$
[9]	EM8502	-
[7][8]	STM32L0	$1\mu A$
[1]	nRF8001, MSP430F2274	$2\mu A + 1\mu A$
[12]	QN9080, BQ22570, AEM30940, HDC2080	$27\mu A$
This work	**PMIC based on AEM10941, RV-3028-C7, MS8892**	$310nA$ (datasheet) $452nA$ (measured)

An autonomous BLE node is presented in [7]. The node uses the BLUENRG-2 BLE SoC as a load, together with an STM32L0-based MCU to control the transition between harvesting and data transmission phase. During the harvesting phase, an AM-1606C solar cell is used to harvest ambient energy down to 200 lux to a $22\mu F$ capacitor. In low-power stop mode, the MCU stays powered-on, operating as a voltage detector to monitor the voltage across the capacitor. As soon as the capacitor voltage reaches a limit of $3.3V$, the node transitions to the data transmission phase. During that phase, the MCU switches to run mode while the BLE IC is utilized in advertising mode to transmit 7-byte advertising data packets with an output power of $+8dBm$ over three channels. The node returns to the harvesting phase based on an interrupt from the low-power timer of the MCU. The quiescent current of the node during its harvesting phase is equal to the operating current of the MCU in stop mode, reported at approximately $1\mu A$. A system with similar characteristics and the addition of the HTS221 IC for moisture and temperature sensing is detailed in [8].

The autonomy of a wearable node for ECG monitoring is detailed in [1]. As loads, the node uses the MSP430F2274 MCU for processing and ADC, the nRF8001 BLE IC for short-range communication, and an analog circuit for ECG measurements, including a pre-amplifier and analog filters. For energy storage, the system uses a $240mAh$ li-ion battery. During operation, the node transitions between: a sleep phase where the load is set into low-power modes; an idle phase where the BLE IC performs advertising; and an active phase where the sensor performs ECG sampling. During sleep mode, nRF8001 is put into idle phase with a reported quiescent current of $2\mu A$ while the MCU consumes approximately $1\mu A$ in low-power mode 3.

A sensor node for the monitoring of physical parameters in reinforced concrete is presented in [12]. The node is powered through wireless power transfer using RF to DC rectifiers to charge a $150\mu F$ capacitor, using two PMICs for energy distribution, BQ22570 and AEM30940. The node transitions between a deep sleep, with an average current of $27\mu A$, and an advertising phase.

To summarize, existing work exhibits a fundamental trade-off between ultra-low quiescent battery-less architectures with limited load capability, and battery-operated systems with higher always-on consumption (Table 1). We improve the state of the

art by integrating a near-zero quiescent drain ($452nA$) on battery-based systems, characteristic of battery-less designs, through a new hardware-orchestrated DPM architecture together with a novel latch circuit for the capacitive sensor.

3 Multi-stage DPM architecture

In order to minimize energy consumption, we have developed a multi-stage DPM architecture that utilizes different operating modes, as depicted in Figure 1: Deep sleep; Wake up; Normal; Overcharge and Shutdown modes. Switching between modes does not involve the use of the MCU to minimize power consumption.

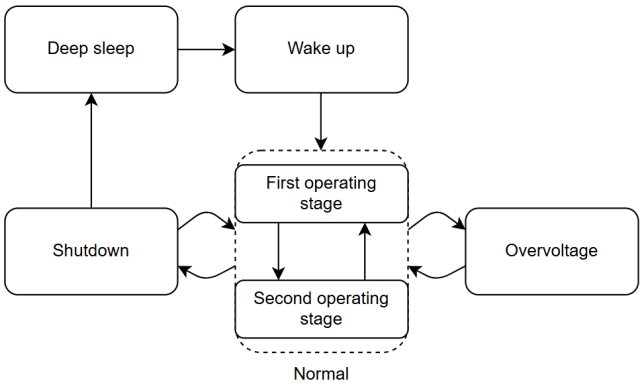

Figure 1: System modes

- In *Deep sleep mode*, the energy storage is deeply discharged and there is no available energy to be harvested from the energy harvester
- When a cold-start voltage of $300mV$ and $2\mu W$ of power become available from the harvester, the PMIC enters *Wake up mode*. If the voltage of the storage element is below a threshold V_{CHRDY}, it will first need to be charged through the PMIC
- When the voltage of the storage element reaches V_{CHRDY}, the PMIC enters *Normal mode*; the outputs Low-Voltage Output (LVOUT) and Always-On Output (AOOUT) are shorted to the buck regulator output and the external DC/DC converter can drive High-Voltage Output (HVOUT). During Normal mode, there can be three different scenarios depending on the voltage of the storage element:
  - If the voltage reaches a higher threshold V_{OVCH}, the PMIC enters *Overcharge mode*; the charging of the storage stops and its voltage is maintained to approximately V_{OVCH} with hysteresis, to prevent damage from overcharging
  - If the storage element's voltage ranges between V_{CHRDY} and V_{OVCH}, the PMIC stays in Normal mode
  - If its voltage drops below V_{CHRDY}, the PMIC enters *Shutdown mode* to prevent deep discharge potentially leading to damage to the energy storage element. If sufficient energy becomes available within approximately $600ms$, the PMIC returns to Normal mode. If not, the circuit returns to Deep sleep mode

The Normal mode can be further divided into two operating stages, the first operating stage and the second operating stage[1]. In the first operating stage, only AOOUT is active. No other components, including the MCU, are powered during this stage. The system is operational and can be woken-up, while maintaining the lowest possible power consumption.

For the system to wake up and enter the second operating stage, the LVOUT power domain must be turned on. The PMIC provides an enable input, 2.2_SW_EN, which controls this power domain. The capacitive sensor and RTC can trigger the enable input of the PMIC in order to turn on the LVOUT switched domain and subsequently power on the system. The capacitive sensor can power the system based on a touch interaction from the user. The RTC on the other hand is programmed during system runtime in order to generate periodical wake up triggers. The enable signal 2.2_SW_EN is latched in the capacitive sensor; it means that the power domain stays on when the enable event is revoked. This feature provides a trigger-style enabling of the $2.2V$ switched domain. An optional, third HVOUT switched domain is also provided by the PMIC, in order to power components that require a higher voltage such as CO_2 sensors. This domain is controlled via the MCU.

A key challenge in the proposed hardware DPM design is translating the momentary triggers from the RTC or capacitive sensor into a persistent *enable* signal that holds the main compute domain *on*. Figure 2 depicts the novel latching mechanism designed to address this problem. This circuit, a primary contribution of this work, is integrated with the capacitive sensor's logic to switch between operating stages.

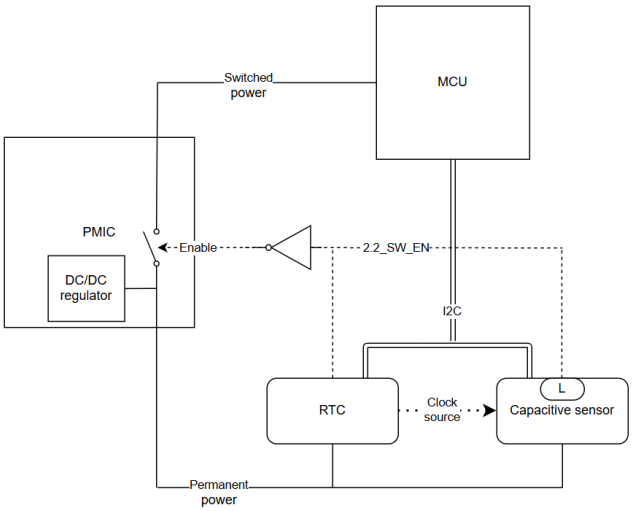

Figure 2: Switching between operating stages

The high level on signal 2.2_SW_EN is achieved by a pull-up resistor. If 2.2_SW_EN is driven to logic low level by MS8892, its internal pull-up resistor is disconnected to avoid the static current flow and further reduce power consumption. The system power state is kept in the internal latch L. The latch L will be set by a touch event or when the RTC pulls the signal 2.2_SW_EN low, thus

[1]Note that the system can be easily extended to include additional operating stages

switching on the switched power via enable. After power-up, the MCU is in control and can read out the wake-up source from the MS8892. It can re-configure the parameters of the MS8892 and the alarm settings of the RTC. The system will return to the deep sleep mode by clearing the latch in the MS8892, which in turn disables the switched power supply. The latch can be cleared by an MCU command via the I^2C bus or alternatively via a digital signal from the MCU to the MS8892 via 2.2_SW_DIS. This design was developed and validated in collaboration with the component manufacturer, Microdul, and was adopted for their commercial MS8892 IC.

Since the interfacing components are always powered, it is critical that they consume minimal energy. For this reason, a single oscillator is permanently running for both the capacitive sensor and the RTC. This is made possible by an external clock input to the capacitive sensor and using the oscillator of the RTC to trigger the periodic capacitive touch detection.

4 Sensor node implementation

To evaluate the DPM architecture presented in section 3, a miniaturized, autonomous SSN was designed and developed, powered by PV energy harvesting. The system integrates an $80MHz$ ARM Cortex-M4 MCU, wireless ICs and multiple sensors, powered by a $10mAh$ battery and a $2cm^2$ high-efficiency PV panel. This section provides details on the SSN implementation, as depicted in Figure 3. A prototype of the printed system can be seen in Figure 4.

The architecture of the sensor node can be divided into four interconnected component blocks, the edge intelligence, power management, wireless and sensing blocks. The edge intelligence block includes the MCU and a non-volatile FRAM for data storage. The power management block integrates the proposed DPM architecture. It includes the PMIC, energy harvester and energy storage, as well as two interfacing components; an RTC and a capacitive touch sensor. The wireless block includes BLE and NFC for short-range communication and LoRa for long-range communication. An additional FRAM module is used for the storage of data from the wireless ICs. Finally, the sensing block includes a CO_2 sensor, accelerometer, temperature sensor, magnetometer and light sensor.

4.1 Edge intelligence block

4.1.1 MCU. For the SSN, an 100-pin MCU based on Cortex-M4 is used, stm32l496vg. The MCU has a frequency of $80MHz$, $1MB$ Flash, $25nA$ quiescent current in shutdown mode and $37\mu A/MHz$ energy consumption in run mode when a Switched Mode Power Supply (SMPS) is used. In this work, stm32l496vg is supplied by an external SMPS, bypassing the internal regulator of the MCU.

The MCU supports two low-power instructions to transition the system into reduced-power states, Wait-For-Event (WFE) and Wait-For-Interrupt (WFI). WFE wakes-up the MCU when the *event* bit in the System Control Block is set, while WFI keeps the MCU in sleep mode until an interrupt occurs. Since the integrated sensors generate hardware interrupts, the system firmware primarily employs WFI, allowing the MCU to enter low-power states between processing cycles and thus minimize overall energy consumption.

The MCU supports eight low-power operating modes, providing different trade-offs between performance and energy use. [13] presents a comparison between the modes in terms of wake-up source, active components and power consumption for an $1.8V$ input voltage. In run mode, the designed SSN supports dynamic voltage scaling to optimize power consumption. The CPU, Flash and SRAM are enabled at a reduced frequency of $2MHz$. In sleep and low-power sleep modes, the CPU halts while peripherals remain active. An interrupt or event can wake-up the CPU. In stop 2 mode, most of the MCU Vcore domain is put into a lower leakage mode, in order to achieve minimal power consumption while retaining SRAM and register contents.

When executing instructions from Flash memory, wait states may be inserted depending on the CPU clock frequency and internal voltage range. The main regulator output voltage, Vcore, supports two modes that affect the total number of wait states and the total power consumption. Range 1, or high-performance range, supports a clock frequency of up to $80MHz$ with a minimum Flash access time for reading. Range 2, or low-power range, has a maximum clock frequency of $26MHz$ and a longer reading time from the Flash memory than Range 1. The correspondence between wait states and CPU clock frequency is presented in [13].

4.1.2 NVM. The proposed DPM architecture of this work is built around the concept of completely switching off most system components during inactive periods to eliminate their quiescent current and thus conserve energy. This also includes the system's memory ICs. For this reason non-volatile memories are used to retain data between shut-down/power up cycles; the status of the running application can be stored in the memory before switching off the device, and can be retrieved on subsequent power ups. For the storage of data between power cycles of this system, the MB85RC64TAPNG-AMEWE1 FRAM is used. The memory performs write operations at the same speed as read operations, supporting a high speed mode of up to $3.4MHz$. The device can be switched off between transfers to conserve energy and be put into sleep mode which reduces current consumption by stooping the internal regulator circuits.

To reduce latency, the system includes two such FRAM chips. The first, which is dedicated to the MCU, is connected via an SPI interface. The second that is dedicated to the BLE IC, communicates with the BLE controller using an I^2C interface.

4.2 Power management block

4.2.1 PV harvester and storage. For the designed SSN, two EXL1-1V20-SM cells were chosen; each cell has a surface of $1cm^2$, yielding a total active surface of $2cm^2$. The cells are connected in parallel to increase the output current of the energy harvester. For energy storage, Powerstream GEB 201212 was used, with a $10mAh$ capacity and dimensions of $2.2 \times 12.5 \times 12.5\text{mm}^3$.

4.2.2 PMIC. For the SSN, we employ a PMIC based on AEM10941, offering a cold-start voltage of $300mV$ and a quiescent current of $200nA$. The PMIC is used for regulating, controlling, and optimizing the power flow from the energy harvester to the battery. It is also used to supply the system's components via three outputs, AOOUT, LVOUT and HVOUT. The inputs and outputs of the PMIC can be seen in Figure 5.

- AOOUT is set around a voltage of $2.2V$. It can continuously drive a small current to supply the capacitive sensor and the RTC even during the sleep mode of the system. The internal

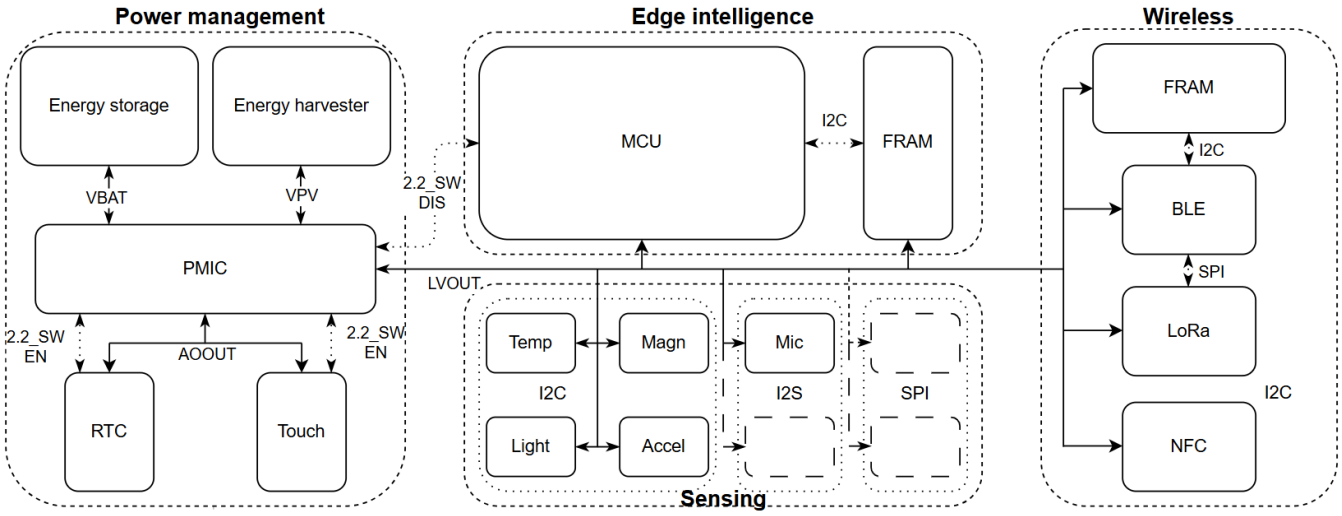

Figure 3: High-level architecture of the SSN

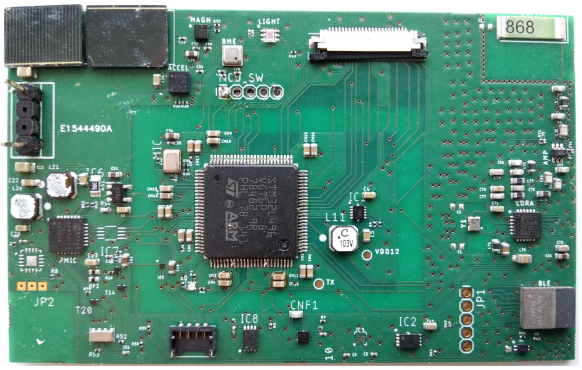

Figure 4: SSN prototype

power gating switch of the PMIC deactivates this output whenever the energy storage is depleted.

- LVOUT is set around a voltage of $2.2V$. It can supply the MCU, sensors or any device able to cope with the $2.2V$ supply, through a power gating switch.
- HVOUT is a $3.3V$ output of the external DC/DC buck converter allowing to supply any components with higher voltage requirements.

4.2.3 RTC and capacitive touch. To minimize standby power, our proposed architecture utilizes an external RV-3028-C7 RTC rather than the MCU-integrated timer. As already discussed, this decoupling allows the MCU and power-demanding peripherals to be switched off during Deep Sleep, while keeping time in the RTC.

The RV-3028-C7 RTC from Micro Crystal is used as part of the power management block and together with the MS8892 capacitive touch sensor are used to wake up the system from the first operating stage. The RTC has a wide operating voltage range of $1.1V$ to $5.5V$ and a low current consumption of typically $45nA$ at $3V$. MS8892 has an energy consumption of $65nA$ when an external clock source is provided.

4.3 Power domains and interfaces

The SSN includes the following domains for the integrated components, corresponding to the outputs of the PMIC. The AOOUT low-leakage $2.2V$ power domain which is always turned on and is used for the proposed power-management architecture, the LVOUT switched $2.2V$ domain and the optional HVOUT $3.3V$ switched power domain that can be used for components that require higher operating voltage, such as CO_2 sensors. For interfacing, I^2C, SPI and I^2S are used. In the I^2C bus, the MCU operates as the master while the RTC, sensors and the NFC operate as slaves. I^2C is also used for the communication between the MCU and BLE with the respective FRAM ICs. The SPI bus is used for the communication between the LoRa module and the BLE controller. Finally, I^2S is used as an interface between the MCU and the microphone.

5 Performance evaluation

This section presents a performance evaluation of the proposed hardware-based multi-stage DPM architecture (section 3), as integrated into the SSN (section 4). We compare the system's performance against a conventional software-based baseline to quantify the energy gains enabled by full hardware power-gating.

To assess the effectiveness of the proposed DPM, the system's quiescent current was experimentally measured in Deep Sleep mode, where only the PMIC's always-on domain, the RTC and the capacitive touch sensor remain active. Measurements were performed using the Nordic Semiconductor Power Profiler Kit II (PPK2), supplying the system at $2.2V$. PPK2 is a standalone unit supporting measurements from $200nA$ up to $1A$, with a resolution between $100nA$ to $1mA$ depending on the range, for a supply voltage of $0.8V$ - $5V$ VCC. The quiescent current of the complete system, including the PMIC's own quiescent current, was measured as $452nA$ at $2.2V$, corresponding to the total energy drain of the system while idle. For

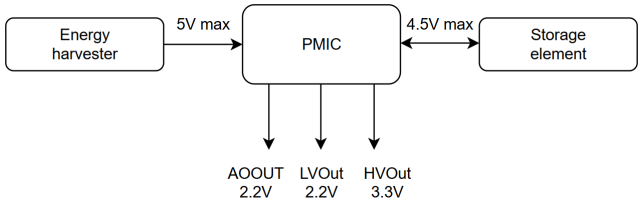

Figure 5: Inputs and outputs of the PMIC

comparison, the theoretical minimum current, estimated from the device datasheets as the sum of the RTC ($45nA$), touch sensor ($65nA$) and PMIC ($200nA$), is approximately $310nA$. The difference between theoretical and experimental values can be attributed mainly to leakage currents in passive components and PCB traces as well as to measurement overhead from the PPK2.

To establish a quantitative baseline, we configured the STM32L4 in Stop 2 mode, representing a state-of-the-art software-based DPM. The measured current consumption for this configuration was approximately $3\mu A$ at $2.2V$. The comparison highlights the significant reduction in idle current achieved through the proposed power gating approach. Specifically, the hardware-based DPM achieves a 6.6x reduction in idle current, effectively eliminating the MCU's quiescent contribution through full domain power-gating.

5.1 Case study: thermal comfort monitoring

To validate practical viability, we deployed the proposed architecture in a wearable thermal comfort monitoring node for indoor environments. The system periodically measures environmental temperature and humidity via BME680, computes thermal comfort metrics on the MCU and informs a room HVAC via BLE ADV frames using OnSemi RSL10. For this case study, the edge intelligence and power management blocks were utilized, as discussed in section 4. The firmware was developed on top of ST's Hardware Abstraction Layer (HAL), as generated by STM32CubeMX.

Assuming a fully charged energy storage, the system starts from the first operating stage. Based on periodic, 10-minute RTC triggers on 2.2_SW_EN, the system proceeds to the second operating stage and the MCU, sensor and BLE ICs are powered. The sensor measures temperature, humidity and atmospheric pressure while the MCU reads previous values from the FRAM. Updated thermal comfort levels are then calculated. The results are transmitted to the HVAC after each measurement in a train of 20 BLE ADV events. In each event, data is advertised in the 3 ADV channels.

Table 2 presents the duration of each operating stage and event, the respective energy requirements in J as well as the energy harvested by the energy harvester cells under 200, 300 and 500 lux. As can be seen, the system spends approximately 10 minutes in the first operating stage, where only the components of the power management block are powered via the 2.2_LOW_LEAK rail.

On an RTC trigger via 2.2_SW, the system transitions to the second operating stage. There, BME680 acquires temperature, humidity and pressure measurements in approximately 1.5 seconds and sends them to the MCU via I2C0. The MCU requires approximately 35 ms to calculate thermal comfort levels. The results are stored in FRAM. Finally, the results are advertised via BLE to the HVAC controller for 2 seconds. The MCU then transitions the system back to the first operating stage. The energy consumption of BME680 for a single measurement of temperature and humidity was measured at $1.1mJ$. The MCU uses an estimated $46.2\mu J$ for the data processing. The total energy consumption of the power management block for the first and second operating stages was measured at $0.6mJ$. The energy consumption for the ADV events is approximately $0.4mJ$. The total energy cost of the system for a cycle of first and second operating stages is approximately $2.14mJ$.

Table 2: Measured energy requirements per event

	Event duration	Energy required
Data acquisition	$1500ms$	$1.1mJ$
MCU	$35ms$	$46.2\mu J$
Always-on domain	10 mins Deep Sleep + $3535ms$ (run mode)	$0.6mJ$
System advertising	$2000ms$	$0.4mJ$
Total per cycle	$3535ms + 10min$	$2.14mJ$

The energy harvested by the EXL1-1V20 cells under 200, 300 and 500 lux for a complete cycle of operation is measured to $26.04mJ$, $38.2mJ$ and $72.42mJ$ respectively resulting in a net energy gain of $23.9mJ$, $36.06mJ$ and $70.28mJ$.

6 Conclusions

This paper addressed the fundamental limitation of quiescent current drain, which renders traditional software-based DPM unfeasible for autonomous systems in energy-scarce, low-light environments. We proposed and implemented a robust, hardware-orchestrated DPM architecture for battery-based sensor nodes, based on [10], that fully power-gates the MCU and all non-essential peripherals, entirely eliminating their quiescent current. A central component of the proposed architecture is the novel hardware latch circuit that enables DPM state transitions, acting as a hardware orchestrator, interfacing with the ultra-low-power RTC and capacitive sensor. This circuit, designed by the authors to provide a persistent enable signal from momentary events for the PMIC , which in turn powers the main compute domain. The novel latch circuit has been already included to the commercial MS8892 IC. Our evaluation demonstrated this architecture achieves a measured sleep-state drain of $452nA$, proving significantly more energy-efficient than traditional software-based sleep modes for battery-based systems.

To validate this architecture, we designed and developed a miniaturized, autonomous SSN integrating the proposed DPM. This system, was validated in a real-world case study at various illuminance levels. It successfully demonstrated that under 200, 300 and 500 lux for a complete cycle of operation, the system had a net energy gain.

Future work will investigate a more granular, multi-stage power-gating approach. While this work gates the entire compute domain, a hybrid model could keep the MCU's SRAM powered in retention mode while gating the core and peripherals. Finally, we will apply this architecture to higher-performance MCUs, such as the STM32U5 and FRDM-MCXN947, both based on the Cortex-M33 core, to quantify the scaling of energy savings as the quiescent current of high-performance cores continues to rise.

References

[1] Klemen Bregar, Tomaž Krištofelc, Matjaž Depolli, Viktor Avbelj, and Aleksandra Rashkovska. 2022. Power Autonomy Estimation of Low-Power Sensor for Long-Term ECG Monitoring. *Sensors* 22, 14 (2022), 5070.

[2] Casper Bröcheler, Thomas Vroom, Derrick Timmermans, Alan van den Akker, Guangzhi Tang, Charalampos S. Kouzinopoulos, and Rico Möckel. 2025. A Segmented Robot Grasping Perception Neural Network for Edge AI. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*.

[3] Sidartha Azevedo Lobo De Carvalho, Daniel Carvalho Da Cunha, and Abel Guilhermino Da Silva-Filho. 2017. Autonomous power management for embedded systems using a non-linear power predictor. In *2017 Euromicro Conference on Digital System Design (DSD)*. IEEE, 22–29.

[4] Vinay Devadas and Hakan Aydin. 2010. On the interplay of voltage/frequency scaling and device power management for frame-based real-time embedded applications. *IEEE Trans. Comput.* 61, 1 (2010), 31–44.

[5] Milad Kokhazadeh, Georgios Keramidas, Vasilios Kelefouras, and Iakovos Stamoulis. 2025. A CNN Compression Methodology for Layer-Wise Rank Selection Considering Inter-Layer Interactions. In *DATE 25 (Design, Automation and Test in Europe)*.

[6] Lars Kröger, Cristian Cioflan, Victor Kartsch, and Luca Benini. 2025. On-Device Federated Continual Learning on RISC-V-based Ultra-Low-Power SoC for Intelligent Nano-Drone Swarms. *arXiv preprint arXiv:2503.17436* (2025).

[7] Roberto La Rosa, Catherine Dehollain, Andreas Burg, Mario Costanza, and Patrizia Livreri. 2021. An energy-autonomous wireless sensor with simultaneous energy harvesting and ambient light sensing. *IEEE Sensors Journal* 21, 12 (2021), 13744–13752.

[8] Roberto La Rosa, Catherine Dehollain, Mario Costanza, Angelo Speciale, Fabio Viola, and Patrizia Livreri. 2022. A Battery-Free Wireless Smart Sensor platform with Bluetooth Low Energy Connectivity for Smart Agriculture. In *2022 IEEE 21st Mediterranean Electrotechnical Conference (MELECON)*. IEEE, 554–558.

[9] Marcel Meli, Manuel Brütsch, Stefan Stajic, Manuel Böbel, David Lorenz, Lukas Hegetschweiler, Dimitris Karanassos, and Charalampos S Kouzinopoulos. 2020. Low light energy autonomous LoRaWAN node. In *2020 IEEE 5th International Symposium on Smart and Wireless Systems within the Conferences on Intelligent Data Acquisition and Advanced Computing Systems (IDAACS-SWS)*. IEEE, 1–6.

[10] Marcel Louis Meli, Sebastien Favre, Benjamin Maij, Stefan Stajic, Manuel Boebel, Philip John Poole, Martin Schellenberg, and Charalampos S Kouzinopoulos. 2023. Energy Autonomous Wireless Sensing Node Working at 5 Lux from a 4 cm2 Solar Cell. *Journal of Low Power Electronics and Applications* 13, 1 (2023), 12.

[11] Gaurav Narang, Raid Ayoub, Michael Kishinevsky, Janardhan Rao Doppa, and Partha Pratim Pande. 2023. Uncertainty-aware online learning for dynamic power management in large manycore systems. In *2023 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*. IEEE, 1–6.

[12] Alassane Sidibe, Gaël Loubet, Alexandru Takacs, and Daniela Dragomirescu. 2022. A multifunctional battery-free bluetooth low energy wireless sensor node remotely powered by electromagnetic wireless power transfer in far-field. *Sensors* 22, 11 (2022), 4054.

[13] STMicroelectronics. 2025. *Ultra-low-power Arm® Cortex®-M4 32-bit MCU+FPU, 100 DMIPS, up to 1 MB Flash, 320 KB SRAM, USB OTG FS, audio, external SMPS.* https://www.st.com/resource/en/datasheet/stm32l496qe.pdf