Intermittent Intelligent Camera with LEO sensor-to-satellite Connectivity

Matteo Nardello
matteo.nardello@unitn.it
University of Trento, Department of Industrial Engineering
Trento, TN, Italy

Luca Caronti
luca.caronti@unitn.it
University of Trento, Department of Industrial Engineering
Trento, TN, Italy

Davide Brunelli
davide.brunelli@unitn.it
University of Trento, Department of Industrial Engineering
Trento, TN, Italy

ABSTRACT

IoT systems can operate efficiently in scenarios with limited or sporadic power availability by utilizing intermittent power sources, such as energy harvesting. Typical implementations are based on fixed-size super-capacitors as energy storage. However, this limits the possible implementations. Large capacitors take more time to charge, resulting in extended off-time following a power failure. Small capacitors charge faster but provide a shorter active time, leading to more frequent power failures. This paper presents a fully intermittent, machine learning-based, low-power smart camera for monitoring applications integrating a two-stage energy harvester. A small first stage supports data acquisition and analysis. The second bigger energy storage is activated only for data streaming to support the integrated sub-GHz Low Earth Orbit transmission radio. This dual-stage energy storage strategy ensures both system reactivity and the capacity to sustain energy-intensive data transmission. The simulation highlights how the fully intermittent pipeline (i.e., for both the neural accelerator and the ARM core) makes the system resilient to power fluctuation and increases the throughput of processed images in ultra-low light conditions by up to 13%.

CCS CONCEPTS

• **Hardware** → **Sensor devices and platforms**; *Sensor devices and platforms*; • **Computer systems organization** → **Sensors and actuators**; • **Computing methodologies** → Machine learning algorithms.

KEYWORDS

Intermittent Computing, LEO Satellite Communication, Intelligent Camera, Energy Harvesting, Power Management

ACM Reference Format:
Matteo Nardello, Luca Caronti, and Davide Brunelli. 2023. Intermittent Intelligent Camera with LEO sensor-to-satellite Connectivity. In *11th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems (ENSsys' 23), November 12, 2023, Istanbul, Turkiye.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3628353.3628550

ENSsys' 23, November 12, 2023, Istanbul, Turkiye

ACM ISBN 979-8-4007-0438-3/23/11.
https://doi.org/10.1145/3628353.3628550

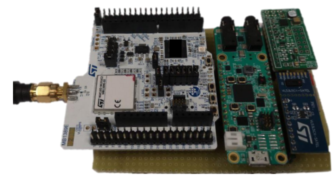

Figure 1: System prototype. From left to right: a Nucleo-WL55JC dev board integrating the LEO radio; MAX78000FTHR Evaluation Board integrating the RGB camera and MAX78000 Neural Accelerator; on the bottom right the ToF sensor and on the top the energy harvester. Not visible in the photo, are the two energy storage elements.

1 INTRODUCTION

Intermittent computing has emerged as a promising approach to developing commercial batteryless IoT devices and addressing their energy constraints. By properly managing intermittent power sources, such as energy harvesting techniques, IoT devices get rid of typical battery problems, like device lifetime, significant environmental impact, and potentially high maintenance costs. Moving away from reliable, fixed-size, energy sources like batteries, poses tight constraints on the device's application, as the system must make judicious use of the available energy. Traditionally an IoT device's operation encompass three different tasks [22]: **1)** Sense; **2)** Compute; **3)** Transmit.

As highlighted by different studies [12], the most energy-hungry task is usually associated with data transmission. Thanks to the recent high-performance, ultra-low-power MCUs, and neural accelerators, it is now possible to implement complex data analysis and filtering directly on the edge with an extra low energy budget. Unfortunately, the supercapacitor-based energy storage must store sufficient energy to complete the most energy-consuming *atomic task*. In this context, an *atomic task* is defined as a task that cannot be interrupted during its execution without data loss. For instance, transmitting a data packet or capturing an image through a visual sensor can be considered "atomic tasks." Typically, the standard practice involves sizing the energy storage capacity to enable the completion of the energy-hungry task.

However, this presents a constraint on potential application scenarios. Larger capacitors require more time to charge, thereby diminishing system reactivity. Conversely, smaller capacitors charge swiftly but offer shorter active durations, resulting in frequent power interruptions or the inability to maintain device functionality. Consequently, researchers have explored the viability of employing dynamic energy storage solutions [13] to construct responsive systems, achieving quick cold starts while simultaneously supporting energy-intensive tasks like powering a transceiver for data transmission. In this paper, we present the development of a battery-free, intermittent visual sensing device for continuous image processing. The camera prototype hosts a Maxim MAX78000[1] Ultra-Low-Power Convolutional Neural Network Accelerator perfectly tailored for tinyML applications. The system incorporates a dual-stage energy harvesting circuit to reduce the cold-start time of the platform while maintaining the capability to support sub-GHz Low Earth Orbit (LEO) radio for data transmission. Leveraging a fully intermittent software pipeline for both the ARM core and the Neural Accelerator, the sensing device has been fine-tuned for energy efficiency and resilience to power fluctuations. The system prototype is presented in Figure 1.

Results show that thanks to the dual-stage energy storage, the visual sensing system can improve application reactiveness to the events, also when the harvestable energy is limited. Moreover, thanks to the LEO radio, the device does not need any additional infrastructure, allowing worldwide data communication. The contributions of this work are:

- The development of capacitor-based dual-stage energy storage to sustain the device's operations under limited energy availability.
- The energy characterization of the sensing platform
- The development of a fully intermittent software pipeline that dynamically selects the suitable energy storage to meet task energy requirements
- Real hardware implementation and software simulation to evaluate the intermittent pipeline.

The paper is organized as follows. Section 2 provides background on energy-harvesting, intermittent computing, and LEO-powered devices. In Section 3, we present an overview of the sensing device hardware, while Section 5 discusses in-depth the intermittent pipeline. Section 5 presents the evaluation and real hardware tests of the platform. Section 6 concludes this work with some final remarks and possible future works.

2 RELATED WORK

Intermittent computing [21] refers to systems and microprocessors that operate on sporadic or unreliable power sources [28]. Instead of relying on a constant and stable power supply, intermittent computing uses irregular energy harvesting from the environment. This could involve capturing energy from ambient light, vibrations, temperature gradients, or even radio waves [34]. The harvested energy is stored in onboard capacitors or batteries, providing short bursts of power for system operation. Intermittent computing enables devices to function in scenarios where access to a continuous power source is impractical or impossible, such as remote locations or harsh environments. Various approaches have been explored to address the challenges posed by power interruptions and system failures. These approaches can be broadly categorized into three types: **i) Full volatile memory checkpointing [24, 31, 35]**. This method incurs significant overhead since it includes unnecessary information, leading to inefficient use of resources; **ii) Custom hardware with transistor-level non-volatility [8, 16, 32, 36]**. This approach offers the advantage

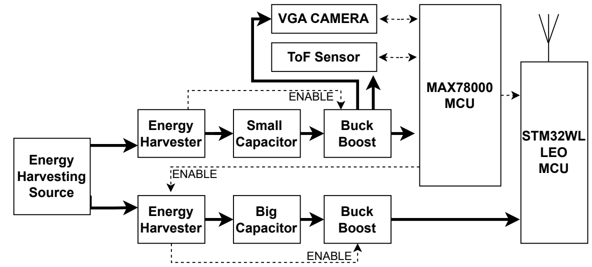

Figure 2: System Overview. From left to right: The energy harvesting source; The two-stage power supply, with the second stage dedicated only to the LEO radio; The MAX78000 neural accelerator along with the sensing devices; and the LEO radio for data streaming. Power traces are in bold, while data lines are dashed.

of preserving data during power interruptions without the need for extensive backup operations, it comes with the drawback of high production costs and the requirement for specialized hardware design; **iii) Task-based backup [3, 23, 25]**. This method selectively backs up only the important and necessary portions of memory, avoiding the overhead of checkpointing unnecessary data. By identifying critical sections and crucial data, this approach optimizes the backup process, reducing resource consumption and improving efficiency. It is also used by Layer-by-Layer Transient Toolchain (LbLTT) [6], which proved to optimize transient workloads; thus, we adopted this framework in our implementation.

As the deployment of IoT sensor networks continues to increase, the volume of sensed information also grows. To avoid transmitting data containing redundant or useless information, researchers have thus investigated Machine Learning-based algorithms to implement efficient data filtering directly on the sensors [1, 12]. In fact, in the case of energy-harvesting devices, the energy demand for communication dominates the total energy budget. Consequently, it makes more sense to transfer ML inference tasks to an edge device or even directly to the sensor [2, 27]. Research has recently proposed methods to improve inference on intermittent computing systems through diverse approaches [10, 20, 30].

Pervasive deployed IoT sensing systems should operate far from communications infrastructure and without the issue of power or battery supply for the infrastructure. Nowadays, the radios in sensor nodes, such as 4G/LTE [11], are expensive per byte, while Bluetooth/Wifi or other IoT pioneering protocols are range-limited [5, 26]. A solution can be represented by chirp spread-spectrum long-range radio technologies technologies [7, 29, 33] but still requires an infrastructure to be built. More recently, researchers have proposed to provide global connectivity exploiting LEO satellites [4, 15]. LEO satellites enable reliable communication services for places with no terrestrial networks, allowing the deployment of IoT devices almost everywhere without worrying about connectivity problems.

3 SYSTEM ARCHITECTURE

The proposed IoT device is a sensing, computing, and communication system including a custom dual-stage power supply, application-level intermittent software pipeline, and a low-power sub-GHz Low

[1]https://www.analog.com/en/products/max78000.html

Earth Orbit transmission radio. Figure 2 shows an overview of the system architecture.

3.1 Sensing Devices

The visual sensing interface is composed of a Time-of-Flight (ToF) sensor and a low-power RGB camera. The ToF sensor, a VL53L5CX from STMicroelectronics, is used as an event generator. It is configured as autonomous low-power mode with an interrupt programmable threshold to wake up the MAX78000 MCU. When an object is detected within a distance less than a custom threshold, an interrupt is generated to wake the MCU up. This permits keeping all the main active parts in sleep mode, avoiding unnecessary energy wastage and making the system more reactive to interesting events. The second sensor is an Omnivision OVM7692 RGB camera already embedded into the MAX78000 evaluation board. Despite the availability of more efficient imaging sensors in the market, we opted to retain this camera for the sake of expediting development and demonstrating the efficacy of our approach. Future iterations will incorporate design refinements utilizing more efficient sensors. The camera has a resolution of 640x480 pixels and is configured to provide QVGA (100x100 pixels) pictures.

3.2 Power Management

Power management is one of the most critical parts of an intermittent visual sensing system. A tandem of requirements must be carefully considered during the design phase: **1)** Responsiveness and reactiveness to the environment. We want to have short recharging periods (i.e., cold starts after power failure) to be able to start processing the information as soon as possible; **2)** Energy Capacity. The energy storage must be tailored to ensure its longest atomic task completes without interruptions; **3)** Efficiency. As the available energy is limited, it is important to lower the quiescent current associated with the selected components, to avoid unnecessary energy wastage. To meet the three design requirements, the proposed implementation boosts a dual-stage energy storage architecture. The first stage, encompassing a small supercapacitor, is responsible for powering the main MCU and the two sensing devices. The second stage is activated by the main MCU only when a data transmission is required and is based on a large supercapacitor, used only to power the LEO MCU Radio. Overall, the energy overhead consumed by the second energy stage is below 1 nA when disabled and less than 1 uA when active. Figure 3 presents the components of the two boosting stages in detail.

3.3 Microcontroller subsystem

At the core of the system lies the MAX78000 device, a high-performance microcontroller unit with advanced neural network capabilities. This MCU integrates an Arm Cortex-M4 Processor with FPU running up to 100MHz, a second 32-bit RISC-V Co-processor running up to 60MHz, and a 64-core CNN accelerator. The two main cores share a total of 512KB Flash and 128KB SRAM, while the neural accelerator has a dedicated 440KB for the weights, 2KB for biases, and 512KB for input data. The accelerator is also connected to the multi-layer bus matrix shared with CPUs, platform memory, and other peripherals. All the internal memory components of MAX78000 are volatile. To enable intermittent operations, an

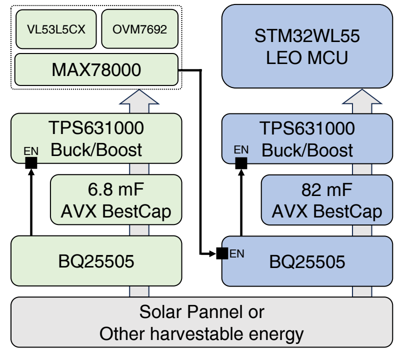

Figure 3: Power management subsystem. On the left the first stage with a smaller supercapacitor tailored for sustaining camera operations (i.e., the most energy-consuming task). On the right, the second stage is activated only when a data transmission is required.

external FRAM (Ferroelectric Random Access Memory) module is connected to the MAX78000 device using SPI protocol. FRAM offers several advantages over traditional memory technologies, such as fast write speeds, non-volatility, and virtually unlimited endurance. This integration allows the system to reliably store critical data, configurations, and intermediate results, ensuring data integrity between power failures and quick access when required.

3.4 Wireless Transceiver

As data transmission front-end, the proposed IoT node exploits a low-cost STM32WL55 MCU, combining an STM32 ultra-low-power microcontroller with a sub-GHz radio transceiver that supports multiple modulation protocols. These include both LoRa, frequency-shift keying modulations, as well as other proprietary sub-GHz standards. The radio is based on a Semtech SX126x transceiver and contains two internal power amplifiers with 15dBm and 22dBm maximum transmit power. In the proposed implementation, the transceiver is configured to generate Kinéis modulation, to enable uplink radio messages transmission directly to Kinéis satellites. Two different modulations are available, namely VLDA4, which allows the transmission of short messages of 3 Bytes, and LDA2 or LDK modulation, which permits to sending standard messages of 19 Bytes (LDK) or 24 Bytes (LDA2). However, due to transmission power limitations of the STM32WL RF front-end, only the VLDA4 modulation can be used with a satisfying probability of receiving a correct message.

3.5 LEO data transmission

Using satellite communication to connect devices has recently been proposed to revolutionize IoT applications. [19] explores the use of such networks for IoT applications where terrestrial alternatives are nonexistent. Thanks to sub-GHz RF that can reach LEO satellites, it's increasingly easier to send small packets to space.

Kineis is a French satellite operator and an IoT connectivity service created in 2018 by CLS, a subsidiary of CNES[2], the French space agency. Kineis satellites use phase shift (BPSK, QPSK, and

[2] https://cnes.fr/en

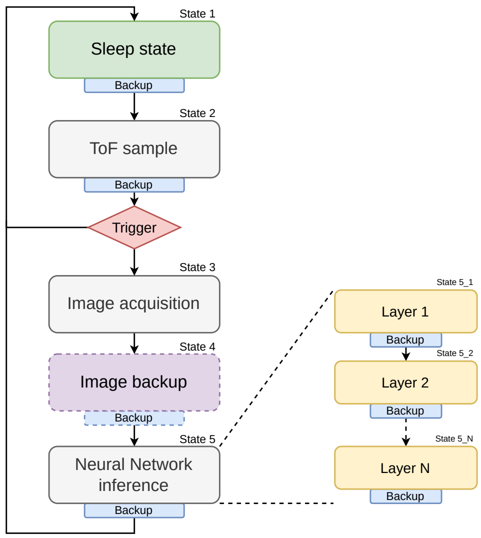

Figure 4: Software pipeline finite state machine. The image backup process and layer-by-layer evaluation are selectively executed based on the remaining energy.

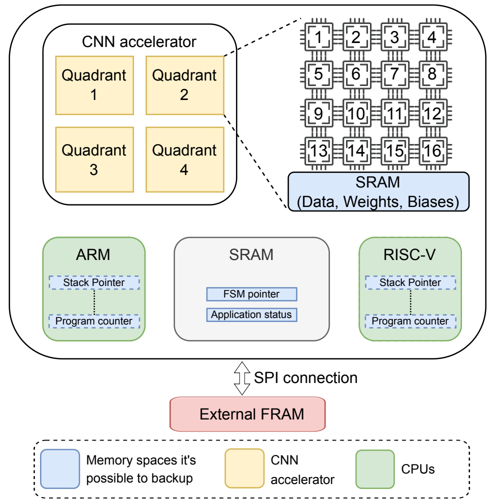

Figure 5: MAX78000 high level architecture overview.

GMSK) modulations adapted to the satellite's communication over different frequency bands (399-401 MHz for environmental sensors and 7 other bands around 400MHz for standard IoT traffic). The standard transmission speed/power is 400 bits/s with at 500mW (27dBm) with a transmission duration between 360ms to 920ms based on the selected modulations, namely VLDA4, LDA2/3, and HD-A3/4 ("VLD" Very Low Data-rate, "LD" stands for Low Data-rate, and "HD" for High Data-rate). A fair-use rule requires 1 minute sleep between two messages (about 1.6% duty-cycle) [18].

The satellites are near the Sun-Synchronous Orbit (SSO), meaning that satellites are passing the same area on Earth, every day at the same time. Communication windows can be easily scheduled by knowing where your device is and the UTC. Kinéis satellites also last longer in space than common CubeSat (eight years vs. a year or two) and have a reentry sequence, so they do not add to the space trash issue once inactive. Currently, as the satellite constellation is not fully completed, when using VDL4 modulation, only two transmission windows per day are available. Usually, with a power of 500mW, one transmission + three repeats have a 99% chance to be received, while at 1W, one transmission (+ two repeats) achieves the same probability [9].

4 INTERMITTENT PIPELINE

The software pipeline is organized as a Finite State Machine (FSM), depicted in Figure 4. It exhibits 5 operational states. The initial state is referred to as *sleep*, characterized by an extremely low energy consumption level. At this stage, the main MCU is kept in sleep mode and everything else is unpowered. Only the ToF sensor is active, waiting for an event. Once an object is detected under a specific threshold, the FSM moves to state 3. State 3 involves the activation of the camera and the acquisition of an RGB image with a resolution of 100x100 pixels. State 4 is exclusively dedicated to the process of saving the captured image in FRAM, ensuring its preservation for further processing. This state is selectively executed based on the remaining energy. If the energy storage is still sufficiently charged, the FSM moves directly into state 5. At state 5, the acquired image is evaluated using a neural network trained for people recognition. To mitigate the problem associated with possible power failures, the machine learning task leverages the Layer-by-Layer Transient Toolchain, LbLTT, as described in [6], enables intermittent execution of individual layers within the neural accelerator, minimizing the risk of data loss. This approach ensures the reliability and integrity of the overall system, reducing the probability of encountering critical failures or losing essential information during the neural network's processing stages. By leveraging the LbLTT technique, the system gains the ability to pause and resume the execution of specific neural network layers, allowing for efficient resource utilization and enhanced fault tolerance.

4.1 MAX78000 memory space

Figure 5 provides an overview of the internal structure of a MAX78000 device and highlights the components that can be backed up in FRAM. Notably, the neural network accelerator consists of four quadrants, each with its dedicated memory space for reading and storing weights or results. These values can be saved in FRAM for persistent storage. Furthermore, the device also features two cores, one based on the ARM Cortex M4 architecture and the other based on RISC-V. If desired, it is possible to directly back up the registers of these processors and restore them at a later stage. However, in the proposed application, we have chosen to save and restore only the state of the FSM directly. By saving the FSM state directly, we ensure that the device can resume its operation precisely from the point of interruption without the need to restore individual processor registers. This approach offers efficiency and simplicity in managing the device's internal state, particularly in the context of the chosen application. The decision to prioritize FSM state preservation over processor register backup stems from the specific requirements and considerations of our application. Preserving the FSM state

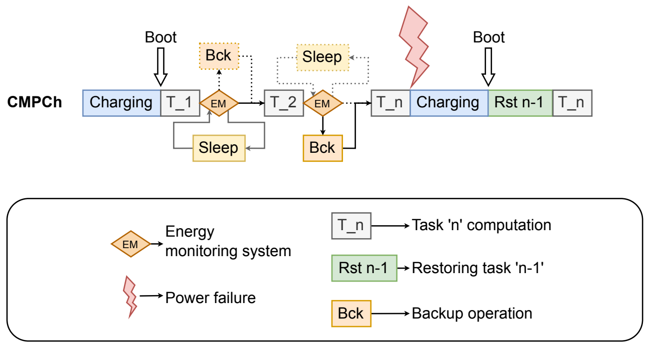

Figure 6: CMPCh Backup policy concept overview.

allows for a seamless continuation of the system's functionality, ensuring the preservation of critical information and minimizing potential disruptions. In contrast, backing up and restoring processor registers might introduce additional complexity and overhead, potentially impacting the overall performance and efficiency of the system.

4.2 Backup policy

Derived from [6], the proposed IoT device implements **CMPCh** (Energy CoMParator CHeckpoint) strategy. Thanks to the MAX78000 integrated low-power voltage comparator, the software pipeline can make informed decisions regarding backup and computation strategy. Depending on the available energy level, the system can choose between three options: **i)** Going into sleep mode to conserve energy until a sufficient amount is accumulated; **iii)** Proceeding directly with the computation; **iii)**Executing a backup of the ongoing task to ensure data integrity. This policy enables efficient resource management and adaptive backup strategies based on the device's energy constraints. This policy optimizes the system's performance while maintaining data reliability by dynamically adjusting backup actions based on energy availability. Figure 6 presents a graphical representation of the proposed backup policy.

5 EVALUATION

We have evaluated the proposed IoT device's power consumption, the impact of having an external SPI FRAM for data retention, the energy associated with data transmission, and the improvement associated with the LbLTT in low energy conditions.

5.1 Power Analysis – MAX78000

Figure 7 presents the energy and time breakdown for the 5 main tasks, namely: Sleep, ToF Sample, Image Acquisition, Image Backup, and Neural Inference. As can be noted, the image acquisition task is one consuming the highest amount of energy. The energy storage of the first stage of the power supply has thus been selected based on the amount of energy needed for image acquisition. Regarding backups, the most prominent data to consider is restoring the image after it has been saved in FRAM. In contrast, the other states merely involve restoring the machine state or other simple registers. However, restoring the 100x100 pixel image takes a considerable amount of time.

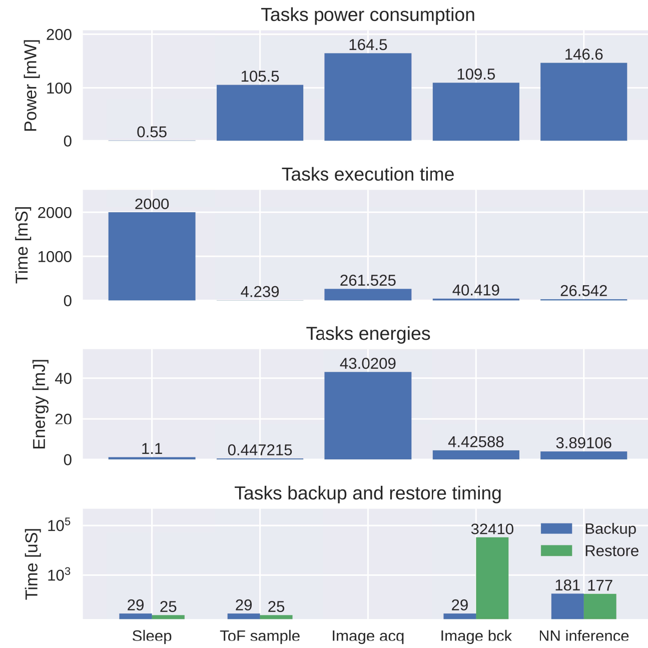

Figure 7: Tasks characterization. From top to bottom: The power consumed by the system during each task; The execution time and the relative energy consumed; The time to perform backup and restore for the various tasks.

5.2 Power Analysis – LEO Transmission

To tailor the energy storage to the energy requirements of the LEO radio, we have collected multiple energy profiles for sending a data packet. The radio is configured to use *VLDA4* modulation, which allows the transmission of short messages of 3 Bytes, enough for transmitting a status indicator. The task's duration from power off to sleep mode after the successful transmission requires 1102 ms, with an average current of 126.93 mA at 3.3 volts. In turn, sending a packet with 3-byte payload requires 0.461 J. Figure 8 presents the current profile while transmitting a data packet.

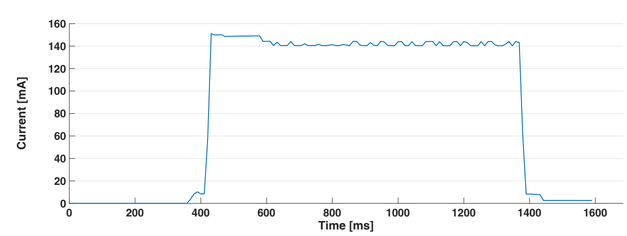

Figure 8: Current profile of the LEO radio during the transmission of a data packet using VLDA4 modulation

5.3 External FRAM vs. theoretical integrated architecture

The decision to utilize an external FRAM module in the hardware setup was primarily driven by the fact that the MAX78000 microcontroller does not incorporate an internal FRAM. However, it is worth

Table 1: Power and time comparison of a theoretical internal FRAM against the used external SPI FRAM.

	SPI FRAM	Internal FRAM*
Time to write/read 1KB	400us	12.5us**
Power consumption	24mW	3mW***
n° of bits wrote in one clock cycle	1	The same as MCU (e.g. 32bit)

* Theoretical evaluation in which FRAM is embedded within the architecture, directly addressable by the CPU, as in the MSP430.
** Assuming a 32-bit architecture.
*** According to [17] at 16MHz and cache hit 100%.

noting that recent advancements in microcontroller technology have led to the development of devices with built-in FRAM capabilities. Therefore, it is plausible that future iterations of the MAX78000 or similar microcontrollers could potentially incorporate internal FRAM for data storage. Using an external FRAM comes with challenges, particularly in timing and power consumption. In this setup, a single-wire SPI connection was employed, meaning that data had to be read or saved serially one bit at SPI clock cycle, introducing a significant overhead in data transfer speed. Additionally, having an external device in the hardware configuration naturally requires additional power consumption. However, in the specific target application developed, the impact of these challenges was relatively minimal. This is because, to backup and restore the state of the FSM, it was sufficient to back up only a single variable. Therefore, the overhead was not a significant concern. Although backup and restoration of the image might pose some power consumption issues, the overall power consumption remained relatively low compared to other tasks in the system. To evaluate the improvement in the case of a MAX78000 with internal FRAM, we interpolated FRAM access in the MSP430FR5994 microcontroller [17] at 1 MHz to match the operation frequency of the MAX78000 ARM core (i.e., 100 MHz). The measurements are presented in Table 1.

5.4 Ultra-Low-Energy LbLTT Improvement Evaluation

To evaluate the computational improvements achievable thanks to LbLTT [6] in ultra-low energy budgets, we have simulated the achievable computational throughput when the harvestable energy is limited. The power trace provided by [14] was used, which provides information on solar irradiation measured in $\frac{uW}{cm^2}$, allowing users to select an appropriate solar panel based on a specific energy constraint. Using the provided power traces allows for evaluating the system's performance under realistic power conditions to assess its behavior once energy is limited. Figure 9 presents the results of this evaluation. Thanks to LbLTT, the application's throughput can improve up to 13%.

6 CONCLUSION

This paper introduces a specialized visual sensing system for scenarios characterized by limited or intermittent power availability.

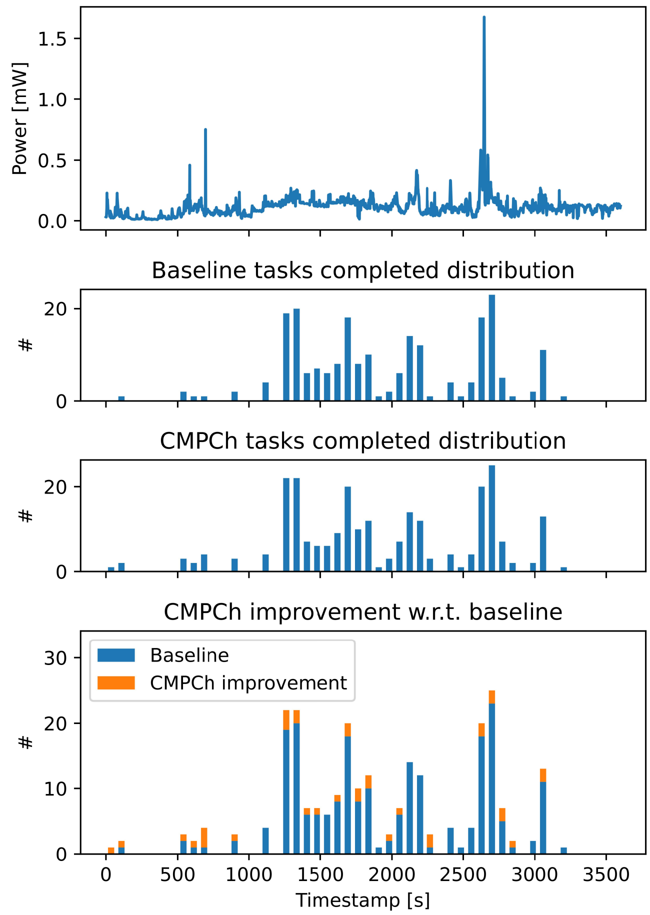

Figure 9: Achievable throughput improvements for the machine learning task thanks to LbLTT under constrained energy budget.

Our device addresses the challenges associated with energy harvesting, cold-start, and responsiveness, using a dynamic energy storage mechanism that intelligently adapts to the specific task.

A small first stage supports data acquisition and analysis. The second bigger energy reservoir is activated only for data streaming to support the integrated sub-GHz Low Earth Orbit transmission radio. Results highlight how the second energy storage introduces negligible energy overhead while allowing a shorter cold start. The paper demonstrates that the fully intermittent pipeline makes the system resilient to power fluctuation and increases the throughput of processed images in ultra-constrained energy budget conditions of up to 13%.

ACKNOWLEDGMENTS

This work was supported by the GEMINI ("Green Machine Learning for the IoT") national research project, funded by the Italian Ministry for University and Research (MUR) by the PRIN 2022 programme (Contract 20223M4HZ4). Moreover, authors would like to thanks Kineis SmartSat program for providing LEO radio software stack and access to the Service portal.

REFERENCES

[1] Andrea Albanese, Matteo Nardello, and Davide Brunelli. 2021. Automated Pest Detection With DNN on the Edge for Precision Agriculture. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 11, 3 (2021), 458–467. https://doi.org/10.1109/JETCAS.2021.3101740

[2] Andrea Albanese, Matteo Nardello, Gianluca Fiacco, and Davide Brunelli. 2023. Tiny Machine Learning for High Accuracy Product Quality Inspection. *IEEE Sensors Journal* 23, 2 (2023), 1575–1583. https://doi.org/10.1109/JSEN.2022.3225227

[3] Abu Bakar, Alexander G. Ross, Kasim Sinan Yildirim, and Josiah Hester. 2021. REHASH: A Flexible, Developer Focused, Heuristic Adaptation Platform for Intermittently Powered Computing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 3, Article 87 (sep 2021), 42 pages. https://doi.org/10.1145/3478077

[4] Mohamed Amine Ben Temim, Guillaume Ferré, and Romain Tajan. 2022. A New LoRa-like Transceiver Suited for LEO Satellite Communications. *Sensors* 22, 5 (2022). https://doi.org/10.3390/s22051830

[5] Davide Brunelli, Massimo Maggiorotti, Luca Benini, and Fabio Luigi Bellifemine. 2008. Analysis of Audio Streaming Capability of Zigbee Networks. In *Wireless Sensor Networks*. Springer Berlin Heidelberg, Berlin, Heidelberg, 189–204.

[6] Luca Caronti, Khakim Akhunov, Matteo Nardello, Kasım Sinan Yıldırım, and Davide Brunelli. 2023. Fine-Grained Hardware Acceleration for Efficient Batteryless Intermittent Inference on the Edge. *ACM Trans. Embed. Comput. Syst.* (jul 2023). https://doi.org/10.1145/3608475 Just Accepted.

[7] Bharat S. Chaudhari, Marco Zennaro, and Suresh Borkar. 2020. LPWAN Technologies: Emerging Application Characteristics, Requirements, and Design Considerations. *Future Internet* 12, 3 (2020). https://doi.org/10.3390/fi12030046

[8] Wei-Hao Chen, Chunmeng Dou, Kai-Xiang Li, Wei-Yu Lin, Pin-Yi Li, Jian-Hao Huang, Jing-Hong Wang, Wei-Chen Wei, Cheng-Xin Xue, Yen-Cheng Chiu, Ya-Chin King, Chorng-Jung Lin, Ren-Shuo Liu, Chih-Cheng Hsieh, Kea-Tiong Tang, J. Joshua Yang, Mon-Shu Ho, and Meng-Fan Chang. 2019. CMOS-integrated memristive non-volatile computing-in-memory for AI edge processors. *Nature Electronics* 2, 9 (01 Sep 2019), 420–428. https://doi.org/10.1038/s41928-019-0288-0

[9] CLS. [n. d.]. *ARGOS Satellite Fleet*. https://www.naos-equipex.fr/content/download/108415/file/10_NAOS_RA1_CLS.pdf (2023, Set 08).

[10] Harsh Desai, Matteo Nardello, Davide Brunelli, and Brandon Lucia. 2022. Camaroptera: A Long-Range Image Sensor with Local Inference for Remote Sensing Applications. *ACM Trans. Embed. Comput. Syst.* 21, 3, Article 32 (may 2022), 25 pages. https://doi.org/10.1145/3510850

[11] Flavio Di Nuzzo, Davide Brunelli, Tommaso Polonelli, and Luca Benini. 2021. Structural Health Monitoring System With Narrowband IoT and MEMS Sensors. *IEEE Sensors Journal* 21, 14 (2021), 16371–16380. https://doi.org/10.1109/JSEN.2021.3075093

[12] Graham Gobieski, Brandon Lucia, and Nathan Beckmann. 2019. Intelligence Beyond the Edge: Inference on Intermittent Embedded Systems. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems* (Providence, RI, USA) *(ASPLOS '19)*. Association for Computing Machinery, New York, NY, USA, 199–213. https://doi.org/10.1145/3297858.3304011

[13] Vasiliki Gogolou, Stylianos Siskos, and Theodore Laopoulos. 2019. A reconfigurable storage unit for autonomously powered IoT nodes. In *2019 Panhellenic Conference on Electronics & Telecommunications (PACET)*. 1–4. https://doi.org/10.1109/PACET48583.2019.8956272

[14] Maria Gorlatova, Michael Zapas, Enlin Xu, Matthias Bahlke, Ioannis (John) Kymissis, and Gil Zussman. 2011. CRAWDAD dataset columbia/enhants (v. 2011-04-07). Downloaded from https://crawdad.org/columbia/enhants/20110407. https://doi.org/10.15783/C77P47

[15] Huawei Huang, Song Guo, Weifa Liang, Kun Wang, and Albert Y. Zomaya. 2019. Green Data-Collection From Geo-Distributed IoT Networks Through Low-Earth-Orbit Satellites. *IEEE Transactions on Green Communications and Networking* 3, 3 (2019), 806–816. https://doi.org/10.1109/TGCN.2019.2909140

[16] Je-Min Hung, Xueqing Li, Juejian Wu, and Meng-Fan Chang. 2020. Challenges and Trends inDeveloping Nonvolatile Memory-Enabled Computing Chips for Intelligent Edge Devices. *IEEE Transactions on Electron Devices* 67, 4 (2020), 1444–1453. https://doi.org/10.1109/TED.2020.2976115

[17] Texas Instruments. [n. d.]. MSP430FR2433 Mixed-Signal Microcontroller. Available at https://www.mouser.com/datasheet/2/405/msp430fr2433-775119.pdf (2022/06/15).

[18] kineis. [n. d.]. *Kinéis satellite operator*. https://connect.kineis.com/discover-satellite-iot/who-is-kineis (2023, Set 08).

[19] Michael Koziol. 2021. *Satellites Can Be a Surprisingly Great Option for IoT*. IEEE Spectrum.

[20] Seulki Lee, Bashima Islam, Yubo Luo, and Shahriar Nirjon. 2020. Intermittent Learning: On-Device Machine Learning on Intermittently Powered System. 3, 4, Article 141 (sep 2020), 30 pages. https://doi.org/10.1145/3369837

[21] Brandon Lucia, Vignesh Balaji, Alexei Colin, Kiwan Maeng, and Emily Ruppel. 2017. Intermittent Computing: Challenges and Opportunities. In *2nd Summit on Advances in Programming Languages (SNAPL 2017) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 71)*, Benjamin S. Lerner, Rastislav Bodík, and Shriram Krishnamurthi (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 8:1–8:14. https://doi.org/10.4230/LIPIcs.SNAPL.2017.8

[22] Dong Ma, Guohao Lan, Mahbub Hassan, Wen Hu, and Sajal K. Das. 2020. Sensing, Computing, and Communications for Energy Harvesting IoTs: A Survey. *IEEE Communications Surveys & Tutorials* 22, 2 (2020), 1222–1250. https://doi.org/10.1109/COMST.2019.2962526

[23] Kiwan Maeng, Alexei Colin, and Brandon Lucia. 2017. Alpaca: Intermittent Execution without Checkpoints. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 96 (oct 2017), 30 pages. https://doi.org/10.1145/3133920

[24] Kiwan Maeng and Brandon Lucia. 2018. Adaptive Dynamic Checkpointing for Safe Efficient Intermittent Computing. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. USENIX Association, Carlsbad, CA, 129–144. https://www.usenix.org/conference/osdi18/presentation/maeng

[25] Kiwan Maeng and Brandon Lucia. 2019. Supporting Peripherals in Intermittent Systems with Just-in-Time Checkpoints. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Phoenix, AZ, USA) *(PLDI 2019)*. Association for Computing Machinery, New York, NY, USA, 1101–1116. https://doi.org/10.1145/3314221.3314613

[26] Kais Mekki, Eddy Bajic, Frederic Chaxel, and Fernand Meyer. 2019. A comparative study of LPWAN technologies for large-scale IoT deployment. *ICT Express* 5, 1 (2019), 1–7. https://doi.org/10.1016/j.icte.2017.12.005

[27] Matteo Nardello, Harsh Desai, Davide Brunelli, and Brandon Lucia. 2019. Camaroptera: A Batteryless Long-Range Remote Visual Sensing System. In *Proceedings of the 7th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems* (New York, NY, USA) *(ENSsys'19)*. Association for Computing Machinery, New York, NY, USA, 8–14. https://doi.org/10.1145/3362053.3363491

[28] Sebin Shaji Philip, Roberto Passerone, Kasim Sinan Yildirim, and Davide Brunelli. 2023. Intermittent Computing Emulation of Ultralow-Power Processors: Evaluation of Backup Strategies for RISC-V. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 42, 1 (2023), 82–94. https://doi.org/10.1109/TCAD.2022.3169108

[29] Tommaso Polonelli, Davide Brunelli, and Luca Benini. 2018. Slotted ALOHA Overlay on LoRaWAN - A Distributed Synchronization Approach. In *2018 IEEE 16th International Conference on Embedded and Ubiquitous Computing (EUC)*. 129–132. https://doi.org/10.1109/EUC.2018.00026

[30] Salonik Resch, S. Karen Khatamifard, Zamshed I. Chowdhury, Masoud Zabihi, Zhengyang Zhao, Husrev Cilasun, Jian-Ping Wang, Sachin S. Sapatnekar, and Ulya R. Karpuzcu. 2022. Energy-Efficient and Reliable Inference in Nonvolatile Memory under Extreme Operating Conditions. *ACM Trans. Embed. Comput. Syst.* 21, 5, Article 57 (dec 2022), 36 pages. https://doi.org/10.1145/3520130

[31] Alberto Rodriguez Arreola, Domenico Balsamo, Anup K. Das, Alex S. Weddell, Davide Brunelli, Bashim M. Al-Hashimi, and Geoff V. Merrett. 2015. 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* (Seoul, South Korea) *(ENSsys '15)*. Association for Computing Machinery, New York, NY, USA, 3–8. https://doi.org/10.1145/2820645.2820652

[32] Simone Ruffini, Luca Caronti, Kasım Sinan Yıldırım, and Davide Brunelli. 2022. NORM: An FPGA-Based Non-Volatile Memory Emulation Framework for Intermittent Computing. *J. Emerg. Technol. Comput. Syst.* 18, 4, Article 73 (oct 2022), 18 pages. https://doi.org/10.1145/3517812

[33] Luca Tessaro, Cristiano Raffaldi, Maurizio Rossi, and Davide Brunelli. 2018. Lightweight Synchronization Algorithm with Self-Calibration for Industrial LORA Sensor Networks. In *2018 Workshop on Metrology for Industry 4.0 and IoT*. 259–263. https://doi.org/10.1109/METROI4.2018.8428309

[34] Sumanth Umesh and Sparsh Mittal. 2021. A survey of techniques for intermittent computing. *Journal of Systems Architecture* 112 (2021), 101859. https://doi.org/10.1016/j.sysarc.2020.101859

[35] Tongda Wu, Kaisheng Ma, Jingtong Hu, Jason Xue, Jinyang Li, Xin Shi, Huazhong Yang, and Yongpan Liu. 2023. Reliable and Efficient Parallel Checkpointing Framework for Nonvolatile Processor With Concurrent Peripherals. *IEEE Transactions on Circuits and Systems I: Regular Papers* 70, 1 (2023), 228–240. https://doi.org/10.1109/TCSI.2022.3208523

[36] Yue Zha, Etienne Nowak, and Jing Li. 2020. Liquid Silicon: A Nonvolatile Fully Programmable Processing-in-Memory Processor With Monolithically Integrated ReRAM. *IEEE Journal of Solid-State Circuits* 55, 4 (2020), 908–919. https://doi.org/10.1109/JSSC.2019.2963005