Low power event detection on microcontrollers

An Empirical Evaluation and Hierarchical Sensing Pipeline

Zion Corbanie
Supervisor(s): Marco Zuñiga Zamalloa, Hao Liu
EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering

June 21, 2026

Name of the student: Zion Corbanie
Final project course: CSE3000 Research Project
Thesis committee: Marco Zuñiga Zamalloa, Hao Liu, Jana Weber

Abstract

Embedded sensing systems relying on energy harvesting — such as electromagnetic radiation, thermoelectric energy, and kinetic energy — generally are not able to harvest sufficient power to function under normal operations for most devices, and thus operate under severe power constraints. To ensure sustainable, battery-free functionality, the microcontroller (MCU) must remain at a low power deep-sleep state during idle periods. It is woken up by a sensor, sending an external hardware interrupt when an environmental event occurs. However, there is a trade-off between a sensor's power consumption, detection range, accuracy, and latency. This paper presents two primary contributions: 1) An empirical evaluation of various sensor wake-up systems. 2) The design and implementation of a multi-stage hierarchical event detection pipeline. This pipeline consists of an ultra-low-power coarse sensor that activates a high accuracy, but higher power sensor, minimizing the current draw while staying reliable.

1. Introduction

Microcontrollers are used all over the world in IoT systems to perform continuous event monitoring, ranging from gesture classification [1] to face detection systems [2]. Traditionally, these devices rely on chemical batteries. However, the continuous replacement of these batteries cause economical constraints and environmental concerns due to toxic electronic waste [3]. Consequently, some research focuses on self sustaining battery-free systems powered by environmental energy harvesting, such as electromagnetic radiation, thermoelectric energy, and kinetic energy [4]. Because the current methods of harvesting energy are not sufficient for most systems, severe restrictions need to be placed on the power consumption of these devices.

Constantly running high accuracy sensors and digital signal processing algorithms consume too much power to meet the strict power requirements. To mitigate this, the microcontroller unit (MCU) is put into a low-power deep-sleep state, consuming on the order of microwatts. The device is woken up by an external hardware interrupt upon event detection.

In 2020, Giordano et al. [2] achieved asynchronous event gating using a VL53L1X Time-of-Flight (ToF) distance sensor to trigger a camera module. While effective, optimizing such a system requires considering the multidimensional trade-off between power consumption, operational range, sensing latency, and accuracy. Generally, more precise sensors with more fine-grained resolutions consume more power.

To address these trade-offs, this paper evaluates such low-power wake-up systems. The primary contribution of this work is the design, implementation, and empirical analysis of a hierarchical event-detection pipeline. This pipeline consists of two sensors in a cascading configuration where an ultra-low-power coarse sensor continuously monitors the environment. When a possible event is detected the high accuracy, but higher power, sensor is triggered. This sensor validates the event before triggering the wake-up of the MCU.

The remainder of this paper is structured as follows. Section 2 establishes the sensor classification profiles and explores wake-up topologies. Section 3 details the empirical measuring methodology and pipeline layout. Section 4 presents quantitative profiling results and evaluates the hierarchical framework. Section 5 discusses performance of the system, and design limitations. Section 6 goes over the environmental and societal effects of this research.

2. Background

Event-detection systems rely on measuring environmental variations and translating these raw measurements into electronic signals. Selecting a sensor requires considering its performance metrics and the nature of the event to be detected. For the facial recognition example in the introduction, sensors are needed that can detect if an object is nearby and in the right position, or within a certain distance range. Consider for example the following trade-off: an ultrasonic distance sensor provides a large measuring range ($\sim 400cm$) but consumes significantly more power ($72.8mW$) compared to a passive infrared (PIR) sensor ($261\mu W$). However, the PIR sensor lacks in distance granularity.

2.1. Hardware

To measure this trade-off under identical hardware conditions, this paper profiles five commercial sensors:

- **Ultrasonic distance sensing:** The HC-SR04 ultrasonic sensor.
- **Proximity sensing:** The TMD27713 proximity sensor.
- **Passive Infrared (PIR) Motion Sensing**: The HC-SR505 motion sensing module.
- **Optical Time-of-Flight (ToF) distance Sensing:** The VL53L1X and VL53L0X laser distance sensors.

The empirical evaluations are conducted using an Arduino Nano 33 BLE as MCU, monitored by the Power Profiler Kit II, a high-resolution power profiler.

2.2. Wake-Up Strategies

To conserve energy, the microcontroller is typically in deep-sleep. This is a state where most of the MCU's features are disabled but can still be woken up by an external, or internal interrupt. There are a couple of strategies considered to wake-up the MCU:

Interrupt wake-up:
Sensors capable of generating an interrupt signal are directly connected to the MCU. The ToF sensors can be configured to trigger an interrupt when a certain condition is met [2]. The TMD27713 proximity sensor, like the ToF sensors, has a programmable interrupt. The HC-SR505 PIR motion sensor outputs a digital 3.3V signal when motion is detected and requires no configuration. All of these sensors have a pin that can be connected to a hardware interrupt pin on the Arduino nano, waking up the device.

Duty cycled polling:
In this method the MCU periodically wakes up from deep sleep using a hardware implemented timer interrupt in the Arduino, and measures the sensor upon waking up. Depending on the measured value, the MCU will decide to go back to sleep or handle the event. The HC-SR04 ultrasonic distance sensor uses a trigger to start ranging, and returns an echo signal where the length of the digital signal is directly proportional to the distance of the measured object. This requires a timer on the MCU to measure the pulse duration digitally.[1]

2.3. Knowledge Gap

While the individual specifications of distance and motion sensors are well-documented, there is a lack of empirical data on the sensors in a full system. Specifically considering the power overhead and latency from detecting an event to waking up from deep-sleep to actually being able to process the event. Such measurements are also seldom tested under identical hardware conditions. This research addresses that gap by profiling five different sensors and evaluating the viability of a hierarchical event-detection pipeline.

3. Methodology & Pipeline Architecture

3.1. Power and Latency Profiling Setup

To accurately capture system's power consumption and latency, a power profiler serves as the power source for the sensor. The profiler provides a stable specified DC voltage while measuring the used current draw over time. For the power measurements, it is not possible to measure the entire system (including MCU) as some sensors require 5V and other require 3.3V. As the Arduino can only run at 3.3V, this would require a power conversion which significantly increases the current draw for sensors at 5V, skewing the results. The power supply of the MCU in deep sleep stays constant during all measurements and is therefore not relevant to the measurement. The MCU power consumption is taken into account for the ultrasonic distance sensor as it is an integral part of the measurement.

To accurately measure the latency, two logic channels are used on the profiler. These channels are synchronous with the power measurements and can be used to measure timing data.

- **Channel 0 (Event assertion)**: Connected to the interrupt pin, marking the exact moment when the sensor validates an event.
- **Channel 1 (Host ready)** Connected to a general purpose IO (GPIO) pin of the MCU. This pin is pulled high when the MCU has woken up and is ready to handle the event, having first disabled the sensor.

Figure 1 shows the general schematic of the measurement setup using a sensor configured by I²C. This schematic differs for the HC-SR505 PIR motion sensor, which is not configured, and the HC-SR04 which does not trigger an interrupt and is read out by the MCU.

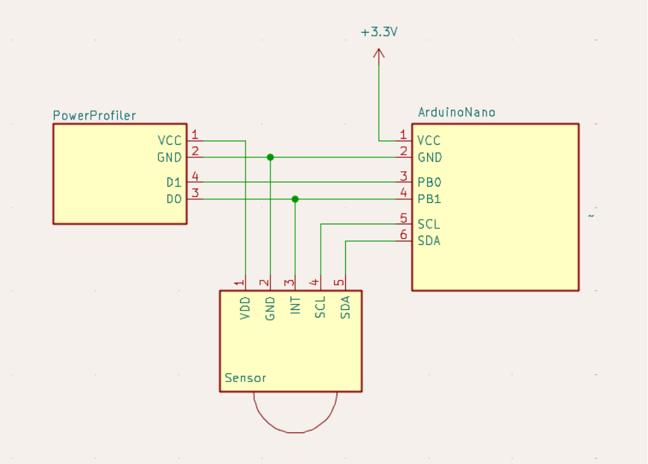

Figure 1: Example schematic example of single sensor profiling setup depicting a sensor, configured by i2c

3.2. Measurement Procedure

For each experiment the following measurements are conducted:

- **Deep-Sleep Power (P_{sleep})**: Continuous current draw when the sensor is in its lowest power state.

---

[1] While a specialized circuit utilizing a RC-integrator to measure the pulse length in a passive analog way can automate this process, this research evaluates the overhead introduced by the MCU polling.

- **Active State Energy (P_{active})**: Total integrated power consumed during sensor excitation, ranging execution.
- **Wake-up Latency ($\Delta T_{\text{latency}}$)**: The temporal delta from the sensor interrupt assertion to the MCU active logic high state.
- **Maximum Effective Range ($R_{\max}$)**: The physical limits of reliable object detection.

Each measurement is conducted indoors with a white A4 sheet of paper as the target. The distance is measured using a tape measure, starting at an out of range distance and bringing the target closer until the event is detected. The time of the event is defined as the moment the sensor is powered on from its sleep state. Since the HC-SR505 PIR motion sensor is not duty cycled, and due to the unknown motion threshold calibration, the start time for this sensor is defined as the moment the event assertion signal goes high. For the TMD27713, which is normally continuous, it was possible to enable a duty cycled mode, which was used for the latency measurement.

Each experiment is conducted five times to ensure consistency. The average of these results are considered. Power consumption is measured using the power profiler.

3.3. Hierarchical Pipeline Design

The main contribution of this paper is a multi-stage event detection pipeline. Instead of continuously running a high accuracy, high power sensor, a multi stage event detection and verification pipeline is established.

With the gathered data from the isolated sensor measurements, a staged pipeline is set up. **Stage A** consists of a coarser but less-consuming sensor (the HC-SR505 PIR), running continuously is used to detect the possibility of an event. In **stage B** a second more precise sensor (the VL53L1X ToF) starts ranging and verifies that the event is indeed worth measuring. Leveraging the cheap energy cost and low latency of the PIR sensor and precision of the ToF sensor ensures the system is accurate and consumes a little power as possible. Furthermore, the output of the PIR motion sensor is connected to the hardware standby pin (`XSHUT`) of ToF sensor, disabling the ToF automatically when there is no motion detected for a certain amount of time.

Because the differing operating voltages present in the system, a 5V to 3.3V DC buck converter is required to convert the voltages. A buck converter was chosen because it is more efficient than the voltage regulator normally present in the Arduino nano which has a quiescent current of nearly 10mA.

A third digital input is used on the power profiler to differentiate between the active stages while measuring the system.

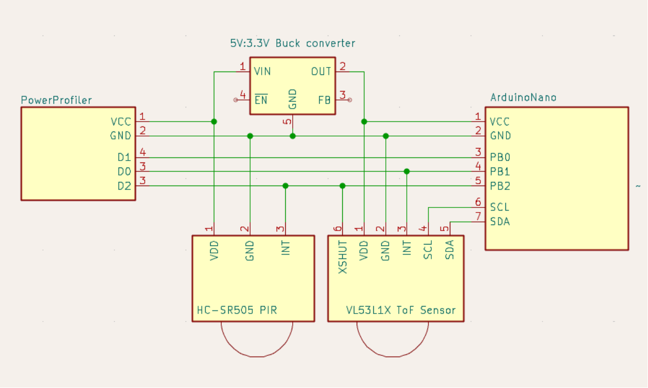

Figure 2: Schematic of the hierarchical pipeline

4. Experimental Results

4.1. Sensor profiling

The following data was gathered from the measuring setup with the power profiler.

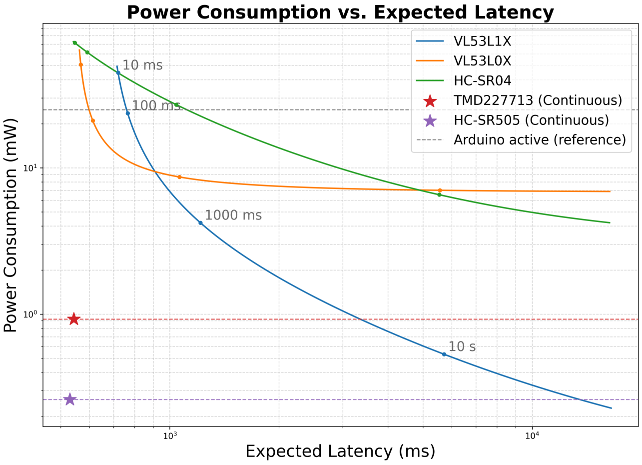

Figure 3: Average power consumption of sensor over the expected latency. Continuous sensors are shown as a point with power consumption reference line. Markers on curves indicate the time between measurements at points 10ms, 100ms, 1s, and 10s

Figure 3 shows the trade-off of power consumption to expected latency. This graph takes into account the active power, sleep power, measure time, and sleep time of each sensor. The proximity and PIR motion sensors are shown as points as their expected latency is not variable. The average power consumption and expected latency are modeled according to the following formulas.

$$P_{\text{sensing}} = \frac{T_{\text{between}} \times P_{\text{idle}} + T_{\text{measure}} \times P_{\text{measure}}}{T_{\text{measure}} + T_{\text{between}}}$$

$$E(T_{\text{latency}}) = \frac{T_{\text{between}}}{2} + T_{\text{measure}} + T_{\text{wake}}$$

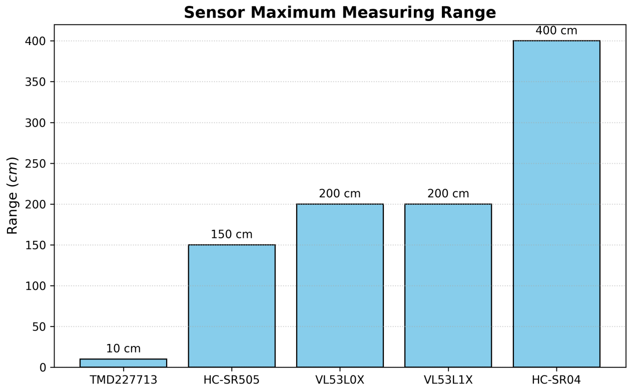

Figure 4: Maximum range range measured in ambient indoor lighting conditions. Rounded to the nearest 50cm with exception for the TMD227713 due to its shorter range.

The maximum usable range measured is shown in Figure 4. Here the measurement result was rounded to the nearest 50cm due to limitations in the accuracy of the measurement setup. The TMD sensor was measured more precisely due to its significantly shorter range of 10cm.

4.2. Individual Sensor Profiles

This subsection considers each sensor with the previous results in order of increasing active power consumption.

The most efficient sensor (in terms of power consumption) is the HC-SR505 motion sensor, but it can not differentiate between close and far objects, nor can it precisely tell where the event is happening. This sensor is great when you want low power, latency, and medium range, but accuracy is not as important.

The next most efficient is the TMD227713 proximity sensor. It can be run in both continuous mode and duty cycled mode but only works with ranges up to 10cm.

After that there are the ToF sensors with similar results. The VL53L1X has a significantly lower power consumption in idle the sleep state, but the VL53L0X has a shorter latency. Their ranges are similar. These sensors also allow you to set a minimum distance, creating a detection window, which differentiates them from the other sensors.

Lastly the HC-SR04 ultrasonic sensor shows the most power consumption but the greatest range.

4.3. Pipeline Implementation & Evaluation

Based on the data from Section 4.1, the HC-SR505 motion sensor and VL53L0X ToF sensor were chosen as sensors for the pipeline. The motion sensor is used in stage A to roughly detect if anything is happening nearby. The ToF sensor is then powered on in stage B to do the precise measurements. Because the motion sensor requires a 5V power supply, and the Arduino and ToF sensor require 3.3V a step-down buck converter was placed in the system such that the system can be measured with the profiler. Figure 5 shows the timing diagram of the system.

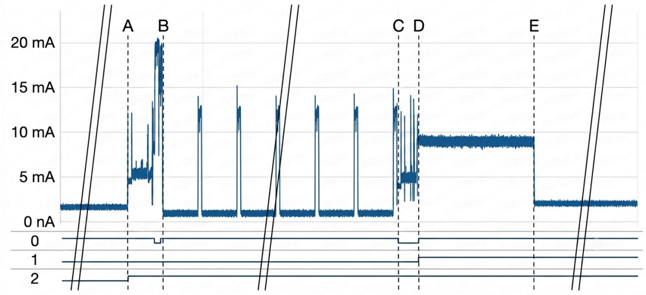

Figure 5: Timing diagram of sensing pipeline. Digital channel D0 is the ToF interrupt output. D1 is the MCU ready response on a successful measurement. D2 is the motion sensor output. Event A shows the detection of an event by sensor A, waking up the MCU to configure sensor B. Event B is when the MCU goes back to sleep and sensor B starts measuring. Event C shows the detection of a valid event powering on the MCU. Event D is the earliest possible response from the MCU, signifying when it can start handling the event. Event E shows the reset to stage A after handling the event.

In stage A, the power consumption is 8.05mW. When the motion sensor detects something it quickly wakes up the MCU to configure the ToF and goes back to sleep to enter stage B. The system here consumes an average of 4.47mW when the sensor is sleeping and 70.1mW while actively ranging. The output of the motion sensor is also connected to the VL53L1X's xshut pin. When this pin is low, the ToF sensor goes into hardware standby, effectively ending the sensing when no more movement happens for a certain time. When the ToF detects a valid event, the MCU is woken up once more to handle the event.

The latency of this system is equal to the VL53L1X's latency, assuming the motion sensor activates before it's considered a valid event.

The range of this system is equal to the minimum of the two sensors, which is 2m.

5. Discussion

5.1. Performance and Metric Trade-offs

The empirical datasets compiled in Section 4 reveal a highly coupled, multi-dimensional trade-off space across power, latency, and range metrics. For instance, while the HC-SR04 ultrasonic sensor provides the greatest maximal range (400cm), its active power consumption (72mW) makes it unsustainable for continuous use in a batteryless system. However, if the latency window of the event of interest is really large — such as a parked car — suddenly the sensor might become viable as seen in Figure 3. Conversely, the HC-SR505 PIR sensor yields a minimal power baseline but operates blindly without spatial depth tracking or object localization features.

The hierarchical event-detection pipeline successfully navigates this trade-off space by decoupling continuous tracking from high-overhead verification. By assigning the ultra-low-power HC-SR505 PIR motion sensor to Stage A, the system maintains a low baseline current draw (8.05mW) during long monitoring phases. The high-overhead VL53L1X ToF sensor is restricted strictly to verified temporal windows (Stage B), meaning its 70.1mA active ranging cost is minimized to the small window where the possibility of a valid event exists.

5.2. Measurement limitations

A comparison between measured range metrics and official datasheet specifications reveals significant differences, as shown in Table 1.

Sensor Line	Datasheet Max Range	Measured Max Range	Performance Deviation
HC-SR04	400 cm	400 cm	0%
HC-SR505	300 cm	150 cm	−50%
VL53L1X	360 cm	200 cm	−45%
VL53L0X	200 cm	200 cm	0%
TMD27713	10 cm	10 cm	0%

Table 1: Measured sensor range deviations compared to datasheet specifications.

The VL53L1X ToF sensor shows a 45% decrease in maximum distance. The datasheet mentions this maximum measurement was taken in the dark. This implies that the difference is likely due to the increased light levels during the measurements from this work. Due to equipment and time limitations during the project phase, a luxmeter was not available to further investigate the effect of the ambient light on the sensor. The other ToF sensor (VL53L0X) showed no deviation from the datasheet, however. Looking at the datasheet, their measurement was taken in indoor lighting conditions as opposed to the dark conditions present with the VL53L1X.

The HC-SR505 PIR motion sensor also showed a significant decrease in the maximum range. Since there is not a lot of information about the sensor, this could be due to various reasons. It is not well defined in the specification of the sensor what constitutes as “motion”. There is also very little information on the calibration of this sensor and the conditions of the test that resulted in the 3m value. The target used in this work was a sheet of A4 paper, which does not have a large heat signature. This could result in the target being harder to detect for the sensor.

5.3. Pipeline limitations

From the results in Section 4.3 it is clear that the pipeline uses considerably more power in stage A and B than expected from the isolated sensor measurements. This is due to a considerable flaw in the system, namely that the motion sensor requires a 5V input and the rest of the system runs at 3.3V, requiring a dc/dc power converter from 5V to 3.3V. In the pipeline a very inefficient buck converter was used with a quiescent current of $787\mu A$ at $5V$, making the largest contribution to the power consumption of the system. There are various viable voltage regulators readily available with quiescent currents into the nano amperes. A better component selection could have drastically improved the results though this remains to be tested.

6. Responsible Research

The recent increase in Internet of Things (IoT) devices brings significant technological advancements, but it also introduces environmental and societal concerns. This section discusses the impact of deploying low-power, energy-harvesting sensing systems in the context of sustainability and technological accessibility.

6.1. Environmental Impact and E-Waste Reduction

Currently, the majority of remote embedded systems rely on traditional chemical batteries. As the number of deployed IoT devices increases, the continuous depletion and required replacement of these batteries generate large quantities of electronic waste (e-waste). The e-waste produced by chemical batteries contains toxic chemicals and heavy metals that cause severe environmental hazards if not disposed of correctly [3]. Furthermore, the manufacturing, distribution, and recycling processes of these batteries contribute heavily to global carbon emissions.

By transitioning toward battery-free systems powered by environmental energy harvesting—such as electromagnetic radiation, thermoelectric energy, and kinetic energy—the ecological footprint of these devices can be drastically reduced. The hierarchical sensing pipeline developed in this thesis contributes to the transition to battery-free systems. By minimizing the power required for continuous event detection, microcontrollers can operate reliably on the power provided by renewable harvesting sources.

Whilst reducing the amount of battery e-waste is positive, no devices have an unlimited lifespan. Therefore facilitating the deployment of IoT devices also creates more e-waste as more devices are being deployed.

6.2. Societal Impact and Technological Accessibility

Beyond environmental conservation, reducing the power constraints on embedded systems democratizes access to smart technologies, making it possible for more people to participate in this field. Systems that do not require frequent maintenance, such as battery swaps, are significantly cheaper to operate over their lifespan. This reduction in overhead makes it economically viable to deploy these IoT devices.

6.3. LLM contributions

In this research LLM's have helped with some coding, in programming the microcontroller, and using python to generate some figures in this work. LLM's have also been used to generate feedback on this report, also helping reword some passages to form a more coherent paper.

7. Conclusion and Future Work

7.1. Conclusion

This paper presented an empirical evaluation of power consumption, latency, and range of five sensors, considering the trade-off between accuracy and power efficiency in sensing systems where power consumption is a limiting factor. Based on these results, a hierarchical sensing pipeline was built consisting of two sensors: a HC-SR505 motion sensor as a coarse first measurement, and a high precision VL53L1X ToF sensor.

While the architecture successfully showed that a high power system can stay offline until necessary, some unexpected power overhead in the system showed up in an inefficient buck converter which was used due to the 5V and 3.3V requirements of the different sensors. Despite this component-level overhead, the sensing pipeline proves to be a viable solution to drastically extending the lifespan of a battery-less embedded system.

7.2. Future work

Several areas in the research have been identified where future work could expand on.

7.2.1. Hardware optimizations:

The effectiveness of the system could be greatly improved by using more efficient components. A voltage regulator with a lower quiescent current could improve the overhead introduced by the motion sensor in stage A.

7.2.2. More sensors

While this research already evaluates five common sensors, there's still a lot of sensors yet to be considered. Having all sensors in the pipeline run on the same voltage could also remove the overhead in the current system.

7.2.3. Improved measurement setup

As mentioned in Section 5, the measurement setup contains some flaws. A more consistent and accurate setup could greatly improve the precision of the evaluation.

Bibliography

[1] H. Liu, Q. Wang, and M. Zuniga, "SolarML: Optimizing Sensing and Inference for Solar-Powered TinyML Platforms," in *2025 Design, Automation & Test in Europe Conference (DATE)*, 2025, pp. 1–7. doi: 10.23919/DATE64628.2025.10993050.

[2] M. Giordano, P. Mayer, and M. Magno, "A Battery-Free Long-Range Wireless Smart Camera for Face Detection," in *Proceedings of the 8th International Workshop on Energy Harvesting and Storage Applications*, ACM, 2020, pp. 48–53. doi: 10.1145/3417308.3430273.

[3] E. M. Melchor-Martínez, R. Macias-Garbett, A. Malacara-Becerra, H. M. N. Iqbal, J. E. Sosa-Hernández, and R. Parra-Saldívar, "Environmental impact of emerging contaminants from battery waste: A mini review," *Case Studies in Chemical and Environmental Engineering*, vol. 3, p. 100104, 2021, doi: https://doi.org/10.1016/j.cscee.2021.100104.

[4] L. Mateu and F. Moll, "Review of energy harvesting techniques and applications for microelectronics," in *VLSI Circuits and Systems II*, J. F. Lopez, F. V. Fernandez, J. M. Lopez-Villegas, and J. M. de la Rosa, Eds., SPIE, 2005, pp. 359–373. doi: 10.1117/12.613046.